

Dynamic Hierarchical Reactive Controller Synthesis

Anne-Kathrin Schmuck, Rupak Majumdar

Abstract

In the formal approach to reactive controller synthesis, a symbolic controller for a possibly hybrid system is obtained by algorithmically computing a winning strategy in a two-player game. Such game-solving algorithms scale poorly as the size of the game graph increases. However, in many applications, the game graph has a natural hierarchical structure. In this paper, we propose a modeling formalism and a synthesis algorithm that exploits this hierarchical structure for more scalable synthesis.

We define local games on hierarchical graphs as a modeling formalism which decomposes a large-scale reactive synthesis problem in two dimensions. First, the construction of a hierarchical game graph introduces abstraction layers, where each layer is again a two-player game graph. Second, every such layer is decomposed into multiple local game graphs, each corresponding to a node in the higher level game graph. While local games have the potential to reduce the state space for controller synthesis, they lead to more complex synthesis problems where strategies computed for one local game can impose additional requirements on lower-level local games.

Our second contribution is a procedure to construct a dynamic controller for local game graphs over hierarchies. The controller computes assume-admissible winning strategies that satisfy local specifications in the presence of environment assumptions, and dynamically updates specifications and strategies due to interactions between games at different abstraction layers at each step of the play. We show that our synthesis procedure is sound: the controller constructs a play which satisfies all local specifications. We illustrate our results through an example controlling an autonomous robot in a known, multistory building.

I. INTRODUCTION

Algorithmic reactive synthesis has recently emerged as a robust methodology to design correct-by-construction controller for specifications given in temporal logics (see, e.g., Girard and Pappas, 2009; Tabuada, 2009; Kloetzer and Belta, 2008; Wolff et al., 2013; Wong et al., 2013). In this technique, one solves a two-player discrete-time game on a graph between the *system* and the *environment* players, where the winning condition is specified in linear-time temporal logic. The game graph is usually obtained as a discrete abstraction of the underlying, possibly continuous or hybrid, dynamics. A winning strategy for the system player in such a game can be computed by algorithmic techniques from reactive synthesis (Zielonka, 1998; Emerson and Jutla, 1991). Such a system winning strategy gives a discrete controller, which can usually be refined to a continuous controller using primitives from continuous control. This controller synthesis methodology has been implemented in symbolic tools (Wongpiromsarn et al., 2011; Mazo et al., 2010; Finucane et al., 2010) and was successfully applied in a number of case studies, e.g., by Wong et al. (2013); Wongpiromsarn et al. (2010).

The two major concerns in the application of reactive synthesis to large problems is (i) the poor scalability of the symbolic game solving algorithms with increasing size of the game graph, and (ii) the limited existence of winning strategies against adversarial environment players in realistic settings. In this paper, we address these challenges by extending the scope of reactive synthesis for control by (i) introducing *local game graphs over hierarchies* as a new decomposed model, (ii) formalizing *hierarchical reactive games* over such models, and (iii) proposing a sound *reactive controller synthesis algorithm* for such games. This algorithm allows for *dynamic specification changes* and uses the construction of *assume-admissible winning strategies* Brenguier et al. (2015) to explicitly model and use environment assumptions.

a) *Local Game Graphs over Hierarchies*: The modeling formalism introduced in this paper allows to exploit the intrinsic *hierarchy* and *locality* of a given large-scale system. This decomposes the controller synthesis problem into multiple small ones. Here, hierarchy means that the game graph allows for the introduction of abstract layers. Locality means that a state at a higher layer naturally corresponds to a sub-arena of the game graph at the next lower layer which is independent from all the other games at the same layer.

As an example, consider an autonomous robot traversing the floors of a building. The lowest layer of the game graph, the game under consideration in existing reactive synthesis techniques, would consist of states defined by grids giving the location and velocity of the robot in each room and each floor of the building, together with additional predicates, such as the location of obstacles, whether the robot is carrying something, or the open-closed status of each door. However, there is a natural hierarchy of abstractions: at the highest layer, we care only about the floors and may ask the robot to move from one floor to another; in the next layer, we would like to know the specific room it is in and specify which room to go next, and only within the context of a room, we may care about where exactly the robot is and where it has to go next. To model this hierarchy, we introduce a set of layers on top of a game graph, each being a game graph itself, where a state at a higher layer (e.g. a room) corresponds to a sub-arena of the game graph at the next lower layer (i.e., all states located inside this room), modeling locality within the hierarchy.

Such hierarchical and local decompositions are also heuristically applied in robotics. Examples are general modeling frameworks, such as hierarchical task-networks (HTN) (Erol et al., 1995) or Object-Action Complexes (OAC) (Kruger et al., 2009), or particular software architectures for incorporating long term tasks and short time motion planning for robots (Kaelbling and Lozano-Perez, 2011; Srivastava et al., 2014; Stock et al., 2015). One could view our abstraction layers, their interaction, and the system dynamics as an equivalent formalism to model task networks. Our controller synthesis algorithms should also apply to design controllers in these formalisms. To the best of our knowledge, the problem of correct-by-construction synthesis for temporal logic specifications (beyond reachability) in the presence of environment assumptions has not been considered by these other formalisms.

Hierarchical approaches for control exist for other correct-by-construction controller synthesis techniques in the control community, such as supervisory control (e.g., Schmidt et al., 2008), hybrid control (e.g., Raisch and Moor, 2005), or continuous control (e.g., Pappas et al., 2000), but these can usually not handle temporal logic specifications.

In many large-scale projects using reactive controller synthesis, such as autonomous vehicles (Hess et al., 2014; Wongpiromsarn et al., 2012) and autonomous flight control (Koo and Sastry, 2002), similar hierarchical and local decompositions are implicitly and informally performed. However, there is no clear theoretical model connecting “low-layer” reactive control and “higher layer” task planning in their work, which is provided by our approach.

b) Hierarchical Reactive Games: To effectively use the constructed hierarchies of local game graphs for reactive controller synthesis, we assume that the specification is also decomposed into a set of local requirements, each restricted to one sub-arena of a particular layer, together with one “global” game at the highest layer. While such a decomposition is not guaranteed to exist for a given specification, it is usually quite natural to exist for specifications over large scale systems with intrinsic hierarchy and locality. For example, for the robot, one may consider the specifications: (i) a floor-layer task “visit all floors”, (ii) a room-layer task “visit all rooms” for each floor, and (iii) a low layer task “if there is an empty bottle [in the current room], reach it and pick it up” for every room.

Synthesizing winning strategies for local games over hierarchies w.r.t. such sets of local specifications becomes challenging due to the interplay between layers both in a bottom-up and a top-down manner. The top-down interplay results because applying a strategy in a higher layer introduces additional specifications for the lower layer. For example, a requested move from one room to an adjacent one requires the local game in this room to fulfill a reachability specification in addition to its local specification. The bottom-up interplay results from the fact that moves in the lowest layer game correspond to moves in all higher layers which might change the strategy. For example, consider a room with two doors to two different adjacent rooms. The higher layer strategy may initially pick one door to continue. However, if this door gets closed before it was reached in the lower layer game, the higher layer strategy might ask to reach the second door instead. Thus, in each local game, winning objectives are generated *dynamically*, based on the strategy at a higher layer, the local specification for the local game and the current system and environment state in the lowest layer.

Intuitively, such interactive hierarchical games are similar to pushdown and modular games (Walukiewicz, 1996; Alur et al., 2003; De Crescenzo and La Torre, 2013), where the local state and the stack determine which (single) local game is played at a particular time point. In contrast, we always play one local game in every layer simultaneously, where visited states in different layers are projections of one another. Therefore, a move in one layer has to be correlated with the games at all other layers at all time steps, giving the dynamic interaction described above.

Our work also relates naturally to abstraction and refinement techniques in game solving, (e.g., Cousot and Cousot, 1977; Henzinger et al., 2000; Abadi and Lamport, 1991), which map “concrete” game structures with “abstract” ones with more abstract timing, to solve a single game for a global specification using different abstraction layers. In comparison, we propose a hierarchical structure where every system state is refined to a whole new local sub-game, having its own specification. Therefore, the game in the higher layer does only proceed for one step once the lower layer local sub-game is completed. In this sense we are “stitching” together solutions of local games in the lowest layer in a particular way which is determined by higher level games, to obtain a solution to the global game.

c) Dynamical Controller Synthesis: Given the hierarchical reactive games described above, we propose a reactive controller synthesis algorithm to solve such games, which allows for *dynamic specification changes* at each step of the play. Intuitively, the controller solves the dynamically constructed local games online and “stitches” their solutions together following the rules of the hierarchical game. Notice that a strategy computed at a level imposes additional conditions on games at lower levels; thus, we use a dynamic controller synthesis algorithm that updates the strategies as the game progresses.

In principle, any algorithm which calculates a winning strategy for a two-player game can be used as a building block to solve local games (e.g., Zielonka, 1998; Emerson and Jutla, 1991; Kupferman and Vardi, 2001; Ehlers and Finkbeiner, 2011; Kupferman and Weiner, 2012). However, these algorithms calculate winning strategies against *any* environment behavior. In most applications, such as our robot example, the requirement that the system wins against any environment strategy is too strong. For instance, in the robot example it is possible, but very unlikely, that an employee keeps an office door closed forever to prevent the robot to fulfill its task. Therefore, recently, assumptions on the environment behavior, which model “likely” behaviors of the latter, were considered to constrain the synthesis problem (see Bloem et al. (2014) and Brenguier et al. (2015) for a detailed overview of recent results). Intuitively, the constrained synthesis problem then asks if

the system can win provided that the environment only behaves according to its assumptions. One type of strategies solving this problem are assume-admissible winning strategies by Brenguier et al. (2015). As this is the most expressive available technique to deal with environment assumptions known by the authors, we use their synthesis algorithm as a building block in our algorithm.

We prove that, whenever the environment meets its assumptions and all dynamically generated local games have a solution, our dynamical synthesis algorithm generates a winning hierarchical play for a given specification, i.e., the algorithm is sound. If these assumptions do not hold, we show that the play gets stuck but does not violate the specification up to this point.

The dynamic nature of our controller is also similar to the receding horizon strategies proposed by Wongpiromsarn et al. (2012); Vasile and Belta (2014), which translate long term goals into current local reachability specifications. This approach allows for a particular two-layer hierarchy and uses time horizons to decompose the synthesis problem locally. However, the general intrinsic hierarchical and local decomposability of a synthesis problem and the interaction of multiple abstract games is not formally exploited. In our presentation, our control synthesis algorithm solves local games completely; however, we can also use a receding horizon controller for each local game.

This paper was motivated by a systems project to build an end-to-end autonomous robotic telepresence system. For the scale of this model, existing reactive synthesis techniques would not work. However, the overall problem has a natural decomposition captured by our proposed model. While this paper focuses on the theoretical foundations of such a formal model and its reactive controller synthesis, we will discuss the implementation and systems aspects of our technique in a different paper.

II. PRELIMINARIES

In this section we first introduce notation and recall existing results from reactive synthesis. Then we discuss a detailed example to motivate our work.

A. Reactive Synthesis Revisited

d) Notation: For a set W , we denote by W^* , W^+ , and W^ω the set of finite sequences, non-empty finite sequences, and infinite sequences, respectively, over W . We write $W^\infty = W^* \cup W^\omega$. For $w \in W^*$, we write $|w|$ for the length of w ; the length of $w \in W^\omega$ is ∞ . We define $\text{dom}(w) = \{0, \dots, |w| - 1\}$ if $w \in W^*$, and $\text{dom}(w) = \mathbb{N}$ if $w \in W^\omega$. We denote by $\text{dom}^+(w) = \text{dom}(w) \setminus \{0\}$ the positive domain of w . For $k \in \text{dom}(w)$ we write $w(k)$ for the k th symbol of w , $[w] = w(|w| - 1)$ for the last symbol of w , and $w|_{[0,k]}$ for the restriction of w to the domain $[0, k]$. Furthermore, $w \cdot w'$ for $w \in W^*$ and $w' \in W^\infty$ denotes the concatenation of two strings. The *prefix relation* on strings is defined by $w \sqsubseteq w'$ if $\exists w'' \in W^* . w \cdot w'' = w'$. Given a set of strings $\varphi \subseteq W^\infty$, we denote by $\overline{\varphi} = \varphi \cup \{w \in W^* \mid \exists w' \in \varphi . w \sqsubseteq w'\}$ the set of strings in φ and all their *finite* prefixes. Slightly abusing notation, we denote by \overline{w} the set $\{w\}$ of all prefixes of the string $w \in W^\infty$.

e) Two-Player Games: A two-player game graph $G = (X, Y, \delta, \rho)$ between environment and system consists of a set of environment states X , a set of system states Y , an environment transition map $\delta : X \times Y \rightarrow 2^X$, and a system transition map $\rho : X \times Y \rightarrow 2^Y$. We assume G is serial, i.e., δ and ρ map each input to non-empty sets. A sequence $\pi \in (X \times Y)^\infty$ with $\pi(k) = (x(k), y(k))$ for all $k \in \text{dom}(\pi)$ is called a *play* in G if

$$\forall k \in \text{dom}^+(\pi) . \left(\begin{array}{l} x(k) \in \delta(x(k-1), y(k-1)) \\ \wedge y(k) \in \rho(x(k), y(k-1)) \end{array} \right). \quad (1)$$

A play π is *finite* if $|\pi| < \infty$ and *infinite* otherwise. The set of all plays is denoted by \mathcal{G} .

We model a *winning condition* in a two-player game as a set of plays $\varphi \subseteq \mathcal{G}$. This set can be represented in different ways, e.g., by an LTL formula or by an ω -automaton. While our results do not assume a particular representation, the latter will determine the algorithm needed to solve the two-player game.

Given a game graph G , a set of initial strings $\mathcal{I} = (X \times Y)^+ \subseteq \mathcal{G}$ and a winning condition $\varphi \subseteq \mathcal{G}$, the tuple $(G, \mathcal{I}, \varphi)$ is called a *game* on G w.r.t. \mathcal{I} and φ . A play $\pi \in \mathcal{G}$ is *winning* (resp. *possibly winning*) for $(G, \mathcal{I}, \varphi)$ if there exists an $n \in \text{dom}(\pi)$ s.t. $\pi|_{[0,n]} \in \mathcal{I}$ and $\pi \in \varphi$ (resp. $\pi \in \overline{\varphi}$). We denote the set of all winning and possibly winning plays for $(G, \mathcal{I}, \varphi)$ by $\text{WinningPlays}(G, \mathcal{I}, \varphi)$ and $\text{WinningPlays}(G, \mathcal{I}, \overline{\varphi})$, respectively.

f) Strategies: A *system strategy* is a *partial* function $f : (X \times Y)^+ \times X \rightarrow Y$ such that¹ $f(w, x) \in \rho(x, [w]_2)$ for all $(w, x) \in \text{dom}(f)$. An *environment strategy* is a *left total*² function $g : (X \times Y)^+ \rightarrow X$ such that $g(w) \in \delta([w])$ for all $w \in (X \times Y)^+$. We denote the sets of system and environment strategies over G by $\mathcal{S}^s(G)$ and $\mathcal{S}^e(G)$, respectively. A play $\pi \in \mathcal{G}$ with $\pi(k) = (x(k), y(k))$ for all $k \in \mathbb{N}$ is *compliant* with $f \in \mathcal{S}^s(G)$, $g \in \mathcal{S}^e(G)$ and $\mathcal{I} = (X \times Y)^+ \subseteq \mathcal{G}$ if there is an $n \in \text{dom}(\pi)$ such that $\pi|_{[0,n]} \in \mathcal{I}$ and for all $k \in \text{dom}(\pi)$, $k > n$, we have

$$x(k) = g(\pi|_{[0,k-1]}) \quad \text{and} \quad y(k) = f(\pi|_{[0,k-1]}, x(k)). \quad (2)$$

¹Here, we write $[w]_2$ for the second component y of the pair $(x, y) \equiv [w]$.

²Due to the serial assumption on G it is possible to assume left total environment strategies.




Fig. 1. Floor plan of the 5th and 6th floor of a six-story building. Using the depicted coordinates, we denote by q_{ij}^k and r_{ij}^k , respectively, the cell and the room in the i th column and j th row of floor k . Furthermore, s_{ij} , $i < j$ denotes the stair case from floor f^i to floor f^j . The workspace of this building is partitioned into grid cells (bottom), rooms (middle) and floors (top) which serve as abstraction layers $l = 0$ to $l = 2$ as discussed in Sec. II-B. The line of dots depicts a path of the robot from the initial state (light gray) to the final state (dark gray) in every layer. Filled circles denote projected states while non-filled circles denote abstract (but not projected) states, as discussed in Expl. 2-3.

The set of plays compliant with f , g and \mathcal{I} is denoted by $\text{CompliantPlays}(f, g, \mathcal{I})$ and we define $\text{CompliantPlays}(f, \mathcal{I}) := \bigcup_{g \in \mathcal{S}^e(G)} \text{CompliantPlays}(f, g, \mathcal{I})$.

A system strategy $f \in \mathcal{S}^s(G)$ is *winning* for $(G, \mathcal{I}, \varphi)$ against $g \in \mathcal{S}^e(G)$, if

$$\begin{aligned} & \forall \pi \in \text{CompliantPlays}(f, g, \mathcal{I}) . \\ & \quad \exists \xi \in \mathcal{G} . \pi \cdot \xi \in \text{CompliantPlays}(f, g, \mathcal{I}) \cap \text{WinningPlays}(G, \mathcal{I}, \varphi). \end{aligned} \quad (3)$$

The set of winning strategies for $(G, \mathcal{I}, \varphi)$ against $g \in \mathcal{S}^e(G)$ is denoted by $\text{WinningStrategies}(G, \mathcal{I}, \varphi, g)$ and we define $\text{WinningStrategies}(G, \mathcal{I}, \varphi) = \bigcup_{g \in \mathcal{S}^e(G)} \text{WinningStrategies}(G, \mathcal{I}, \varphi, g)$.

A system strategy f is *dominated* by a system strategy f' in the game $(G, \mathcal{I}, \varphi)$ (see Brenguier et al. (2014, Def.3)), if for all $g \in \mathcal{S}^e(G)$ holds

$$f \in \text{WinningStrategies}(G, \mathcal{I}, \varphi, g) \Rightarrow f' \in \text{WinningStrategies}(G, \mathcal{I}, \varphi, g).$$

A system strategy which is not dominated is called *admissible*. The set of admissible strategies in the play $(G, \mathcal{I}, \varphi)$ is denoted by $\text{AdmissibleStrategies}(G, \mathcal{I}, \varphi)$.

g) The Synthesis Problem: The (unconstrained) synthesis problem takes as input a game $(G, \mathcal{I}, \varphi)$ and asks if there is a winning system strategy for the game. In most applications, the requirement that the system wins against any adversarial environment strategy is too stringent. The constrained synthesis problem additionally takes as input an assumption that models “likely” behaviors of the environment as a set of plays $\zeta \subseteq \mathcal{G}$. Intuitively, the constrained synthesis problem asks if the system can win provided that the environment player is restricted to play strategies that ensure ζ . In the presence of environment assumptions, the synthesis problem looks for *assume-admissible winning strategies* for the system (see Brenguier et al. (2015) for a discussion why this is an appropriate notion).

By swapping the roles of system and environment we can equivalently define winning and admissible strategies for the environment in the game (G, \mathcal{I}, ζ) as before. Then a system strategy f is *assume-admissibly winning* for $(G, \mathcal{I}, \varphi)$ w.r.t. ζ (Brenguier et al. (2015), Rule AA) if

$$\begin{aligned} & f \in \text{AdmissibleStrategies}(G, \mathcal{I}, \varphi) \quad \text{and} \\ & \forall g \in \text{AdmissibleStrategies}(G, \mathcal{I}, \zeta) . f \in \text{WinningStrategies}(G, \mathcal{I}, \varphi, g). \end{aligned} \quad (4)$$

It should be noted that every winning strategy is assume-admissibly winning w.r.t. any assumption, but not vice-versa.

B. Example

To illustrate the theoretical results and their accompanying assumptions in this paper, we consider a robot that moves in a six story building with known floor plan, depicted in Fig. 1 (bottom) for floors 5 and 6.

To model this problem as a two-player game graph G , we partition the workspace into small cells which form a uniform grid. The resulting grid cells are enumerated by an index set Q . By assuming that the robot can only be in one grid cell at a time, the system state set is given by $Y = Q$. We furthermore define the set of environment states by $X = 2^Q$, where a state $x \in X$ is a set containing all grid cells which are currently occupied by an obstacle.

This modeling formalism implies that each grid cell in Fig. 1 (bottom) represents a system state. We model additional properties by adding other binary variables. For example, by adding a predicate `Bottle` to the system state, we model

whether the robot is carrying a bottle or not. As this additional variable might be true in any grid cell, the resulting system state set would consist of two copies of the grid world in Fig. 1 (bottom), where one is annotated with `Bottle` and the other one is not. To keep notation simple, such additional predicates are mostly neglected in this example.

The system transition map ρ in G results from applying an appropriate abstraction method for continuous dynamics, e.g., Tabuada (2009), while adding the obvious restrictions that (i) the robot cannot move into an obstacle-occupied cell, and (ii) the robot can only move to adjacent cells that are not separated by a wall. For the environment transition map δ several levels of detail can be used to model the movement and (dis)appearance of obstacles, see e.g., Wong et al. (2013); Vasile and Belta (2014) for examples.

Now consider a task for the robot which asks it to reach a specific room on a specific floor. This corresponds to a *reachability* winning condition. In our setting, the winning condition is captured by the language of all plays π such that there exists $k \geq 0$ with $\pi(k) = (x(k), y(k))$ and $y(k)$ is a cell in the specified room. (It can easily be described in linear temporal logic as well.) The synthesis problem for this specification over the game graph G finds a strategy (a controller for the robot) that ensures that the robot eventually reaches the room.

There are two challenges in applying reactive synthesis in this scenario. First, the requirement that the robot must reach the room against all possible environments is too stringent. In such a robot motion example the environment player naturally has a very rich set of possible moves. For the specification considered above, the environment can simply keep a couple of doors closed forever to prevent the robot to reach its goal. However, this adversarial behavior is very unlikely in a real world application as, e.g., employees in an office building will always eventually visit/exit their office. This is the reason why we introduce environment assumptions that constrain the problem. A natural environment assumption allowing to realize the above specification models that all staircases are always eventually unblocked, all doors get always eventually re-opened, and moving obstacles always eventually allow a passage to exit a room.

As discussed in Brenguier et al. (2014), one cannot simply perform reactive synthesis w.r.t. environment assumptions by considering the implication $\zeta \Rightarrow \varphi$ that requires the controller to ensure φ holds only on plays satisfying ζ . This is because the robot may win the game by simply violating the environment assumption (for example, by blocking a door and preventing the environment from opening it). Thus, we consider assume-admissible strategies in this paper.

The second challenge is that of scalability. In any realistic model of our problem, the number of states is so large that existing reactive synthesis tools do not scale. Our main contribution in this paper is to scale up reactive synthesis techniques by considering *local* structure. We now consider this in more detail.

As depicted in Fig. 1, there is a natural hierarchy on the states of the workspace imposed by rooms and floors. That is, the workspace can also be partitioned using the set of rooms R or the set of floors F as index sets.³ This partition introduces two abstraction layers with decreasing precision with system state sets $Y^1 = R$ and $Y^2 = F$. The set of environment states in layers 1 and 2 are defined as the set of closed doors $X^1 = 2^D$ and the set of blocked staircases $X^2 = 2^S$, respectively. Even though the three layers in Fig. 1 are constructed separately, there is a natural abstraction relation between system states $f \in F$, $r \in R$, and $q \in Q$. A system state q is obviously related to the system state r if the grid cell q is “inside” room r . Furthermore, a door d is marked as `closed` if all cells intersecting with this door are occupied by an obstacle (usually being the door itself in this case), inducing a relation between environment states of layers 0 and 1. In Section III, we present *abstract game graphs* (AGGs) which capture such hierarchies in reactive games.

The abstraction relations naturally decompose every layer in the example into small, local game graphs located “inside” a higher level system state: the game graph G is decomposed in local game graphs G_r , $r \in R$. This is possible for this example as the set of possible moves in one room is independent from the part of the environment state that does not belong to this context, e.g., all the obstacles contained in the set x that are not located inside this room. In Section IV, we introduce *local game graphs* (LGGs) which decompose AGGs to model this locality within the hierarchy.

To exploit this local structure in reactive synthesis, we additionally require that the specification is also given as a set of local specifications, one for each local game; otherwise, there is no obvious way to automatically break a global specification into local synthesis problems. For example, for the reachability task, one can consider a specification of reaching a room at the higher layer, and reaching from one point of a room to a prescribed exit point in the lower layer. Correspondingly, notice that the environment assumptions can also be decomposed into layers.

As a second example, consider the more complex task:

*“Collect all empty bottles in the building and return
them to the kitchen in the 5th floor.”*

This task can be manually decomposed in a natural fashion as follows. The level 2 task asks the robot to visit all floors of the building and to return to floor 5 whenever its capacity to carry empty bottles is reached. While in one floor, the level 1 task asks the robot to visit all rooms until the carrying capacity is reached, and to visit the kitchen whenever the latter is true and the robot is in floor 5. Finally, the level 0 tasks ask the robot to search for empty bottles in a single room, approach each bottle and pick it up. In this paper we assume that both the system specification and the environment assumptions are

³For simplicity we model the stairs as a separate room and always “attach” the downward stairs to the respective floor.

already given in a decomposed manner. The automatic decomposition of a global winning condition into local ones is an orthogonal, difficult, problem.

In Section IV-B, we define *hierarchical reactive games* (HRGs) by combining the set of LLGs over hierarchies with a set of local winning conditions and a set of local environment assumptions. This generates a set of local games over an LGG w.r.t. a local specification φ and a local assumption ζ .

The main challenge for reactive synthesis for HRGs is that the games played at the various layers interact. That is, a strategy at a higher layer (“go to the kitchen”) introduces additional constraints at the lower layer (“the higher level strategy requires that the robot should go to the exit that takes it to the kitchen”). In Section V, we provide a synthesis algorithm that computes a dynamic controller for HRGs. The controller computes assume-admissible strategies for each local game, and dynamically updates the winning conditions and strategies through the hierarchy. We prove that the algorithm is sound and that it aborts the game only when a local subgame cannot be won by the system against admissible strategies of the environment.

III. HIERARCHICAL DECOMPOSITION

We now introduce a hierarchy of L two player game graphs where the higher layers are a more abstract representation of the original game graph at layer $l = 0$.

A. Layering, Abstract Plays, and Timescales

Let $G = (X, Y, \delta, \rho)$ be a game graph. A sequence $\langle X^0, Y^0 \rangle, \langle X^1, Y^1 \rangle, \dots, \langle X^L, Y^L \rangle$ is a *layering* of G if (i) $X^0 = X$ and $Y^0 = Y$, and (ii) for each $l \in [1, L]$, there exist *abstraction functions* $\alpha_s^l : Y^{l-1} \rightarrow Y^l$ and $\alpha_e^l : (X^{l-1} \times Y^{l-1}) \rightarrow X^l$.

Notice that while the system abstraction function maps system states at level $l-1$ to system states at level l , the environment abstraction function α_e^l maps a pair (x, y) of environment and system states at level $l-1$ into an environment state at level l . This allows us to incorporate the loss of direct control with increasing abstraction level, as illustrated in the following example.

Example 1: Consider the robot in Sec. II-B and assume that the system states of layer 0 are extended by the binary variable `Bottle`, resulting in the state $\{q, \text{Bottle}\}$ if the robot is in cell q and carries a bottle and the state $\{q\}$ if the latter is not true. In this example, a transition from state $\{q\}$ to $\{q, \text{Bottle}\}$ is enforceable in layer 0 if there is a bottle in cell q (which can be modeled by a corresponding environment variable) assuming that the robot can always pick up a bottle when it is in this cell

Now assume that the specification in the room level asks the robot to go to the kitchen, if it is carrying a bottle. To realize this task, a strategy in layer 1 does not need to *enforce* the robot to pick up a bottle in a particular room (because it might not actually know in which rooms bottles are located) but only *observe* that the latter happened. This intuition can only be modeled if `Bottle` is included in the environment states rather than the system states of layer 1. To be able to trigger this environment variable in layer 1 when the robot picks up a bottle, the tuple $(x, \{q, \text{Bottle}\}) \in X^0 \times Y^0$ must be projected to an environment state $\{\text{Bottle}\} \cup x' \in X^1$ using the map α_e^1 . \triangleleft

For notational convenience, we define the composition of abstraction functions $\alpha_e^{l\uparrow} : (X \times Y) \rightarrow X^l$ and $\alpha_s^{l\uparrow} : Y \rightarrow Y^l$ as

$$\forall x \in X, y \in Y . \alpha_e^{l\uparrow}(x, y) = \alpha_e^l(\alpha_e^{l-1}(\dots \alpha_e^1(x, y))), \quad (5a)$$

$$\forall y \in Y . \alpha_s^{l\uparrow}(y) = \alpha_s^l(\alpha_s^{l-1}(\dots \alpha_s^1(y))) \quad (5b)$$

and the special cases $x = \alpha_e^{0\uparrow}(x, y)$ and $y = \alpha_s^{0\uparrow}(y)$.

A layering induces an abstraction for a play $\pi \in \mathcal{G}$ for each layer $l > 0$ as follows. Given a game G , a play $\pi \in \mathcal{G}$, and layers $\langle X^l, Y^l \rangle_{l=0}^L$ with abstraction functions α_e^l and α_s^l , we define the set of *abstract plays* $\Pi = \{\pi^l\}_{l=0}^L$ of π by $\pi^l \in (X^l \times Y^l)^\infty$ with $\pi^l(k) = (x^l(k), y^l(k))$ s.t.

$$\forall k \in \text{dom}^+(\pi) . \left(\begin{array}{l} x^l(k) = \alpha_e^{l\uparrow}(x(k), y(k-1)) \\ \wedge y^l(k) = \alpha_s^{l\uparrow}(y(k)) \end{array} \right) \quad (6)$$

and $\pi^l(0) = (\alpha_e^{l\uparrow}(x(0), y(0)), \alpha_s^{l\uparrow}(y(0)))$.

Intuitively, the abstract plays in Π are an abstraction of the play π which becomes coarser the higher the layer, as multiple system and environment states are clustered into one state in a higher level. Specifically, this implies that state changes occur less frequently in a higher level than in the play π as outlined in the following example.

Example 2: Consider the path of the robot depicted by filled cycles in Fig. 1 (bottom). This path represents the system state component y of a play $\pi \in \mathcal{G}$. Applying the second line of (6), this sequence y can be abstracted to layer $l = 1$ and

$l = 2$ as follows.

$$\begin{aligned} y &= q_{22}^5 & q_{23}^5 & q_{33}^5 & q_{43}^5 & q_{53}^5 & q_{54}^5 & q_{55}^5 & q_{56}^5 & \dots \\ y^1 &= r_{11}^5 & r_{11}^5 & r_{11}^5 & r_{21}^5 & r_{21}^5 & r_{21}^5 & r_{22}^5 & r_{22}^5 & \dots \\ y^2 &= f^5 & f^5 & f^5 & f^5 & f^5 & f^5 & f^5 & f^5 & \dots \end{aligned}$$

The abstract sequences y^1 and y^2 are depicted in Fig. 1 (middle) and (top), respectively. The state changes in levels 1 and 2 correspond to changes in rooms and floors, respectively. While the state at level 0 changes in each time step, observe that state transitions in layers 1 and 2 only happen irregularly and not at every time point. It should be noted that environment states in layer 1 and 2, i.e., the set of closed doors and blocked stairs, can change independently from system state changes and is not illustrated in Fig. 1. \triangleleft

Expl. 2 illustrates that an abstract play π^l is usually not turn-based. To obtain a turn-based game and to remove redundant information, we introduce a new time scale for every layer which is triggered by changes in the system states in an abstract game π^l as follows. Given a play $\pi \in \mathcal{G}$ and a layer $l \in [0, L]$, the *timescale transformation* κ^l of π in layer l is the identity function if $l = 0$, and defined by the strictly monotone sequence $\kappa^l \in \mathbb{N}^\infty$ s.t.

$$\kappa^l(0) = 0, \tag{7a}$$

$$\forall m \in \text{dom}(\kappa), m > 0, k \in [\kappa(m-1), \kappa(m)) \text{ . } \uparrow \tag{7b}$$

$$\begin{aligned} &\left\lfloor \bullet \right. y^l(k) = y^l(\kappa^l(m-1)) \neq y^l(\kappa^l(m)) \\ &\text{and } \forall k > \lceil \kappa^l \rceil \text{ . } y^l(k) = y^l(\lceil \kappa^l \rceil), \end{aligned} \tag{7c}$$

otherwise. The set of *projected plays* $\check{\Pi} = \{\check{\pi}^l\}_{l=0}^L$ of π with $\check{\pi}^l = (\check{x}^l, \check{y}^l)$ is defined as the sub-sequence of the abstract play π^l at time points given by κ^l for every $l \in [1, L]$. Formally,

$$\forall k \in \text{dom}(\kappa^l) \text{ . } \check{\pi}^l(k) = \pi^l(\kappa^l(k)). \tag{8}$$

A projected play $\check{\pi}$ is called *infinite* if $|\check{\pi}| = \infty$ and *finite* otherwise. While plays $\pi \in \mathcal{G}$ can always be made infinite (by the serial assumption on the transition relations), its projection $\check{\pi}^l$ to layer $l > 0$ need not be infinite. For example, if the robot from Sec. II-B should just move within room r_{11}^5 , this obviously induces an infinite play π . However, its projection to the room layer is given by $\check{\pi}^1 = r_{11}^5$, i.e., $\check{\pi}^1$ is finite with length 1.

Example 3: Consider the abstract sequences y^1 and y^2 in Expl. 2. Using (7) and (8) their induced time scale transformations are given by

$$\kappa^1 = 0 \ 3 \ 6 \ \dots \quad \text{and} \quad \kappa^2 = 0 \ 20$$

and the resulting projections for layer 1 and 2 are given by

$$\check{y}^1 = r_{11}^5 \ r_{12}^5 \ r_{22}^5 \ \dots \quad \text{and} \quad \check{y}^2 = f^5 \ f^6$$

corresponding to changes in rooms and floors respectively at those times. In Fig. 1, system states of projected plays are depicted by filled circles, whereas states only belonging to abstract plays are depicted by non-filled circles. \triangleleft

It can be easily shown (see Lem. 1 in App.) that the range of the timescale transformation κ^{l+1} is a subset of the range of κ^l ; if there is an event at the $(l+1)$ st layer, there is a corresponding event at the l th (and so, in each lower) layer. Using this observation we can simplify notation by defining

$$\kappa_l^{l+1}(k) := (\kappa^l)^{-1}(\kappa^{l+1}(k)) \tag{9}$$

to denote the position in the l th layer of the k th event in the $(l+1)$ st layer.

B. Abstract Game Graphs

Using the notion of abstract states and plays from the previous section, we now construct game graphs for every layer l . We remark that the actual game is only played in the lowest layer, i.e., in the game graph G , and the higher layers only model projected plays of this game.

Definition 1: Let $G = (X, Y, \delta, \rho)$ be a game graph, and $\langle X^l, Y^l \rangle_{l=0}^L$ a layering of G using the abstraction functions α_e^l and α_s^l . Then we define the set of *abstract game graphs* (AGG) $\{G^l\}_{l=0}^L$ for each layer $l \in [1, L]$ by $G^l := (X^l, Y^l, \delta^l, \rho^l)$ s.t.

$$x' \in \delta^l(x, y) \Leftrightarrow \left(\exists \pi \in \mathcal{G}, y' \in Y^l \text{ . } \left(\begin{aligned} &\pi^l(\kappa^l(0)) = (x, y) \\ &\wedge \exists k \in (0, \kappa^l(1)] \text{ . } \pi^l(k) = (x', y') \end{aligned} \right) \right) \tag{10a}$$

$$y' \in \rho^l(x, y) \Leftrightarrow \left(\exists \pi \in \mathcal{G}, x' \in X^l \text{ . } \left(\begin{aligned} &\pi^l(\kappa^l(1) - 1) = (x', y) \\ &\wedge \pi^l(\kappa^l(1)) = (x, y') \end{aligned} \right) \right) \tag{10b}$$




Fig. 2. Generation of system and environment transitions for layer $l = 1$ from a play π as formalized in Def. 1 and discussed in Expl. 4.

and for $l = 0$ by $G^0 := G$. \triangleleft

Intuitively, the maps δ^l and ρ^l collect all transitions that can occur in projected plays $\tilde{\pi}^l$ of possible lowest level plays $\pi \in \mathcal{G}$, as illustrated in the following example. It should be noted that all lowest level plays π are existentially quantified in (10), i.e., all possible plays in the lowest layer are considered.

Example 4: Consider the play $\pi \in \mathcal{G}$ and its abstract play π^1 depicted in Fig. 2. The existence of the play π introduces the depicted system and environment transitions using (10a) and (10b), respectively. Observe that the construction considers every environment change (induced by the play π) as an environment transition from the environment state at the last triggering instance indicated by κ . Furthermore, system transitions are only generated at triggering times. It can be seen in Fig. 2 that the environment state in layer $l > 0$ possibly changes multiple times before a system state change follows. \triangleleft

The construction in Def. 1 allows us to prove that projected plays $\tilde{\pi}^l$ as defined in (8) are also plays in the game graph G^l , i.e., $\tilde{\pi}^l \in \mathcal{G}^l$. Intuitively, the proof shows that there always exist transitions, as the ones emphasized in Fig. 2, connecting system and environment states at triggering times.

Proposition 1: For any game G , any play $\pi \in \mathcal{G}$, and any $l \in [0, L]$, we have that $\tilde{\pi}^l$ is a play in G^l , i.e., $\tilde{\pi}^l \in \mathcal{G}^l$.

Proof: The claim follows directly from Lem. 2 in App. as (1) holds for $\tilde{\pi}^l$ and G^l when we pick $n = \kappa^l(m + 1)$ in (35). \square

IV. CONTEXT-BASED DECOMPOSITION

A set of AGGs imposes an abstraction hierarchy on top of a given game graph G . However, AGGs by themselves are not enough to decompose a synthesis problem. For example, if the winning condition is given by a set of plays on the lowest layer, the induced abstraction layers cannot be exploited by a synthesis algorithm. In order to derive an efficient synthesis technique, in this section, we introduce the second ingredient: *local* winning conditions, which induce *local game graphs*.

Roughly, a *local* winning condition for the game G^l at layer l is a set of abstract plays π^l whose states belong to a single state at layer $l + 1$. For example, reaching a different floor is a local specification at layer 2. A synthesis procedure to enforce φ^L would require solving games at lower levels; in our example, the robot will have to successively reach a set of rooms, followed by the stairs to achieve its goal. Each of these “lower level” games occur in, roughly, the “local” game structure defined by states in the lower level that map to the current state of the higher level. We formalize this notion as *local game graphs*.

A. Local Game Graphs over Hierarchies

Fix a layer l and consider the games G^l and G^{l+1} . Consider a system state $\nu \in Y^{l+1}$. A first attempt to define a local game is to restrict the game G^l to the set of system states $\{y \in Y^l \mid \alpha_s^{l+1}(y) = \nu\}$. However, this is not sufficient, because plays in the local game should be allowed to leave the region specified by ν for one step at the end. This is necessary to ensure that plays in consecutive local games can be concatenated to form a play over the game graph G^l without formalizing a special reset action, as e.g., used in modular games by Alur et al. (2003). To account for these states, we introduce the Post operation:

$$\text{Post}^l(\nu) := \left\{ \nu' \in Y^l \mid \left(\begin{array}{l} \nu' \neq \nu \\ \wedge \exists x \in X^l . \nu' \in \rho^l(x, \nu) \end{array} \right) \right\}. \quad (11)$$

Including the one-step post states allows us to view the actual game as a layer 0 game and use the hierarchical and local decompositions as modeling formalism for hierarchical controller synthesis only.

Considering environment states instead of system states, a straightforward restriction to a context ν is not naturally given by $\alpha_e^{l+1\uparrow}$, as the following example shows.

Example 5: Consider the example from Sec. II-B and its floor plan depicted in Fig. 3. Recall from Sec. II-B that an environment state $x \in X^0$ contains all grid cells that are occupied by an obstacle. However, by playing a game in room r_{11}^5 one is only interested in obstacles that are located inside $Y_{r_{11}^5}^0$. \triangleleft

Therefore, instead of using $\alpha_e^{l+1\uparrow}$ to restrict X^l to context ν , we use a restricting function τ_ν^l . For Expl. 5, the map $\tau_{r_{11}^5}^1$ simply maps the set x of obstacle locations to the subset $x' \subseteq x$ of such locations that are inside the striped area in layer




Fig. 3. Floor plan from Fig. 1. The striped areas in layers 0 and 1 correspond to $Y_{r_{11}^5}^0$ and $Y_{f^5}^1$, respectively. The three arrows denote context changes requested by layer l which induce a reachability specification for layer $l-1$ whose initial and goal states are depicted in light and dark gray, respectively.

0 of Fig. 3. For notation convenience, we define τ^L as the identity map. Using the above intuition, we define *local game graphs* as follows.

Definition 2: Given an AGG G^l , the *local game graph* (LGG) $G_\nu^l := (X_\nu^l, Y_\nu^l, \delta_\nu^l, \rho_\nu^l)$ at layer l restricted to $\nu \in Y^{l+1}$ consists of

$$X_\nu^l := \{\tau_\nu^l(x) \mid x \in X^l\} \text{ and} \quad (12a)$$

$$Y_\nu^l = Y_{\nu \downarrow}^l \cup Y_{\nu \uparrow}^l \quad (12b)$$

$$\text{s.t. } Y_{\nu \downarrow}^l := \{y \in Y^l \mid \nu = \alpha_s^{l+1}(y)\} \text{ and} \quad (12c)$$

$$Y_{\nu \uparrow}^l := \left\{ y' \in Y_{\nu \uparrow}^l \mid \left(\begin{array}{l} \nu' \in \text{Post}^l(\nu) \\ \wedge \exists y \in Y_{\nu \downarrow}^l, x \in X_\nu^l. y' \in \rho^l(x, y) \end{array} \right) \right\}, \quad (12d)$$

and transition maps $\delta_\nu^l : X_\nu^l \times Y_{\nu \downarrow}^l \rightarrow 2^{X_\nu^l}$ and $\rho_\nu^l : X_\nu^l \times Y_{\nu \uparrow}^l \rightarrow 2^{Y_\nu^l}$ defined as:

$$(x' \in \delta^l(x, y) \wedge y \in Y_{\nu \downarrow}^l) \Rightarrow \tau_\nu^l(x') \in \delta_\nu^l(\tau_\nu^l(x), y) \text{ and} \quad (13a)$$

$$(y' \in \rho^l(x, y) \wedge y \in Y_{\nu \downarrow}^l \wedge y' \in Y_{\nu \uparrow}^l) \Rightarrow y' \in \rho_\nu^l(\tau_\nu^l(x), y). \quad (13b)$$

We write $[\mathbb{G}] := \left\{ \{G_\nu^l\}_{\nu \in Y^{l+1}} \right\}_{l=0}^{L-1} \cup \{G^L\}$ for the set of LGGs over G . \triangleleft

Example 6: Consider the example from Sec. II-B and its floor plan depicted in Fig. 3. The striped areas in layers 0 and 1 correspond to the context restricted system state sets $Y_{r_{11}^5}^0$ and $Y_{f^5}^1$, respectively. It is easy to see that $Y_{r_{11}^5}^0 = \{q_{25}^5, q_{43}^5\}$ and $Y_{f^5}^1 = \{s_{56}^5\}$, while layer $l=2$ is not decomposed. \triangleleft

In the robot example of Sec. II-B the generated set of LGGs is “truly local” in the sense that the local system dynamics do not depend on environment variables from other contexts. E.g., an obstacle in another room r' does not influence the dynamics of the robot in room $r \neq r'$. This inherent decomposability of the system dynamics, similar to the natural relations among states of different layers, is a feature of the system we want to control which is necessary for the subsequently proposed synthesis algorithm and formalized in the following assumption.

Assumption 1: For every layer $l \in [0, L-1]$ and context $\nu \in Y^{l+1}$ it holds for all $x \in X^l$ and $y \in Y_{\nu \downarrow}^l$ that

$$y' \in \rho^l(x, y) \Rightarrow y' \in \rho_\nu^l(\tau_\nu^l(x), y). \quad (14)$$

It should be noted that the right hand side of (14) uses ρ_ν^l instead of ρ^l . Therefore, $\rho_\nu^l \subseteq \rho^l$ if Ass. 1 holds, which implies that in this case (13) holds in both directions.

Similarly to Prop. 1 we can prove that the part of a play π^l that takes place in context ν is actually a play in G_ν^l . However, to formalize this we need to define *local plays* which are projected to the current context. Given a set of LGGs $[\mathbb{G}]$, a play $\pi \in \mathcal{G}^0$ and its sets of abstract and projected plays Π and $\tilde{\Pi}$, the *local restriction* of π^l and $\tilde{\pi}^l$ is defined for all $m \in \text{dom}^+(\tilde{\pi}^l)$ by

$$\pi_\downarrow^l(m) := (x_\downarrow^l(m), y^l(m)) \quad \text{with} \quad x_\downarrow^l(m) := \tau_{y^{l+1}(\kappa^l(m)-1)}^l(x^l(m)) \quad \text{and} \quad (15a)$$

$$\tilde{\pi}_\downarrow^l(m) := (\tilde{x}_\downarrow^l(m), \tilde{y}^l(m)) \quad \text{with} \quad \tilde{x}_\downarrow^l(m) := \tau_{y^{l+1}(\kappa^l(m)-1)}^l(\tilde{x}^l(m)). \quad (15b)$$

The restriction of $x^l(m)$ (resp. $\tilde{x}^l(m)$) at time $k = \kappa^l(m)$ is defined w.r.t. the last system state $y^{l+1}(k-1)$ as $y^{l+1}(k)$ is only available after the next system move that is depended on $x(k)$. The local restriction $\tilde{\pi}_\downarrow^l$ of the projected play introduces

a sequence \check{p}_\downarrow^l of local projected plays defined by

$$\forall m \in \text{dom}^+(\check{\pi}^{l+1}) . \check{p}_\downarrow^l(m-1) := \check{\pi}_\downarrow^l|_{[\kappa_i^{l+1}(m-1), \kappa_i^{l+1}(m)]} \quad (16a)$$

$$\text{and } \check{p}_\downarrow^l(\text{end}(\check{\pi}^{l+1})) = \lceil \check{p}_\downarrow^l \rceil := \check{\pi}_\downarrow^l|_{[\lceil \kappa_i^{l+1} \rceil, \text{end}(\check{\pi}^l)]}, \quad (16b)$$

where $\text{end}(w) = |w| - 1$ denotes the time of the last element of w . We write $[\check{p}]_\pi := \left\{ \check{p}_\downarrow^l \right\}_{l=0}^{L-1} \cup \{ \check{p}_\downarrow^L \}$ for the set of all such sequences induced by π , where $\check{p}_\downarrow^L(0) = \check{\pi}^L$ and $\text{end}(\check{p}_\downarrow^L) = 0$.

Example 7: Consider the play π whose y -component is depicted by filled cycles in Fig. 1 (bottom). For illustration purposes, assume a static environment with a closed door between room r_{11}^5 and r_{12}^5 , denoted by the binary variable d , and an obstacle in q_{63}^5 . The closed door, which is an environment variable for layer 1, corresponds to obstacles in q_{24}^5 and q_{25}^5 for layer 0. For this play, the local plays contained in the set $[\check{p}]_\pi$ are given by the following strings.

$$\begin{aligned} \check{p}_\downarrow^0(0) &= (\{q_{24}^5, q_{25}^5\}, q_{22}^5)(\{q_{24}^5, q_{25}^5\}, q_{23}^5)(\{q_{24}^5, q_{25}^5\}, q_{33}^5)(\{q_{24}^5, q_{25}^5\}, q_{43}^5) \\ \check{p}_\downarrow^0(1) &= (\{q_{24}^5, q_{25}^5\}, q_{43}^5)(\{q_{63}^5\}, q_{53}^5)(\{q_{63}^5\}, q_{54}^5)(\{q_{63}^5\}, q_{55}^5) \\ &\vdots \\ \check{p}_\downarrow^0(7) &= (\{\perp\}, q_{62}^6)(\{\perp\}, q_{63}^6) \\ \check{p}_\downarrow^1(0) &= (\{d\}, r_{11}^5)(\{d\}, r_{21}^5)(\{d\}, r_{22}^5)(\{d\}, r_{32}^5)(\{d\}, s_{56}) \\ \check{p}_\downarrow^1(1) &= (\{d\}, s_{56})(\{\perp\}, r_{12}^6)(\{\perp\}, r_{11}^6)(\{\perp\}, r_{21}^6) \\ \check{p}_\downarrow^2(0) &= (\{\perp\}, f^5)(\{\perp\}, f^6). \end{aligned}$$

where $\{\perp\}$ denotes that no obstacles are present. Due to the definition of Y_ν^l in Def. 2, contexts of neighboring cells overlap. This is also visible by the above local plays, which overlap for one time instant. E.g., the state $(\{q_{24}^5, q_{25}^5\}, q_{43}^5)$ belongs both to $\check{p}_\downarrow^0(0)$ and $\check{p}_\downarrow^0(1)$, which are the local plays in context $Y_{r_{11}^5}^0$ and $Y_{r_{21}^5}^0$, respectively. As we use the convention that the environment moves first, the environment variables of such overlapping states are always restricted to the context, which is currently left. \triangleleft

Proposition 2: Let $[\mathbb{G}]$ be a set of LGGs and \mathcal{G}_y^l the set of plays in G_y^l . Furthermore, let $\pi \in \mathcal{G}$ and $[\check{p}]_\pi$ its induced set of local projected play sequences. Then it holds for all $l \in [0, L-1]$ and $m \in \text{dom}(\check{\pi}^{l+1})$ that

$$\check{p}_\downarrow^l(m) \in \mathcal{G}_{\check{y}^{l+1}(m)}^l. \quad (17)$$

Proof: (17) follows by combining the last lines of (36a) and (36b) in Lem. 3 proven in App. . \square

B. Hierarchical Reactive Games over Sets of LGGs

We have seen in the example of Sec. II-B that the motivation for constructing LGGs comes from the natural decomposability of system dynamics, environment assumptions and tasks into local and global components which are naturally restricted to a context $\nu \in Y^{l+1}$. Recall that local specifications should intuitively only contain finite strings to eventually allow progress in the higher layer upon completion of the local task. This observation is formalized as follows. Given a set $[\mathbb{G}]$ of LGGs, layer $l \in [0, L-1]$, and context $\nu \in Y^{l+1}$, the sets

$$\varphi_\nu^l \subseteq (X_\nu^l \times Y_{\nu_l}^l)^* \cap \mathcal{G}_\nu^l \quad \text{and} \quad \zeta_\nu^l \subseteq (X_\nu^l \times Y_{\nu_l}^l)^\infty \cap \mathcal{G}_\nu^l \quad (18)$$

are the *local system specification* and the *local environment assumption* for G_ν^l , respectively. The sets $\varphi^L \subseteq \mathcal{G}^L$ and $\zeta^L \subseteq \mathcal{G}^L$ are a system specification and an environment assumption for G^L , respectively. We define sets of local system specifications and local environment assumptions over $[\mathbb{G}]$ as

$$[\varphi] := \left\{ \left\{ \varphi_\nu^l \right\}_{\nu \in Y^{l+1}} \right\}_{l=0}^{L-1} \cup \{ \varphi^L \} \quad \text{and} \quad [\zeta] := \left\{ \left\{ \zeta_\nu^l \right\}_{\nu \in Y^{l+1}} \right\}_{l=0}^{L-1} \cup \{ \zeta^L \}. \quad (19)$$

A winning strategy for a local specification in layer $l+1$ induces transitions from a state (x, y) to a (possibly different) state (x, y') . As $y, y' \in Y^{l+1}$ are different contexts for layer l , this order of contexts must be obeyed by the strategy in layer l . Therefore, we need a proper translation of transitions in level $l+1$ into reachability specification for local games in layer l and combine these specifications with the given low level tasks. Formally, the reachability specification for a layer $l \in [0, L-1]$ in context $\nu \in Y^{l+1}$ w.r.t. the next context $\nu' \in \text{Post}^{l+1}(\nu)$ is defined by

$$\psi_\nu^l(\nu') := \begin{cases} \{w \in (X_\nu^l \times Y_\nu^l)^* \cap \mathcal{G}_\nu^l \mid [w] \in Y_{\nu_l}^l \cap Y_{\nu'_l}^l\}, & \nu \neq \nu' \\ \{(X_\nu^l \times Y_{\nu_l}^l)^\omega \cap \mathcal{G}_\nu^l\}, & \nu = \nu' \end{cases} \quad (20)$$

and the combination of $\psi_\nu^l(\nu')$ with a local task $\varphi_\nu^l \in [\varphi]$ is defined by

$$\phi_\nu^l(\nu') := \{ \xi \cdot \xi' \mid \xi \in \varphi_\nu^l \wedge [\xi] \cdot \xi' \in \psi_\nu^l(\nu') \}. \quad (21)$$

Example 8: Consider the floor plan in Fig. 3 and assume that the robot is in state q_{22}^5 corresponding to the states r_{11}^5 and f^5 in layers $l = 1$ and $l = 2$, respectively, as indicated by the light gray coloring. Now assume that the controller in layer $l = 2$ requests a context change from f^5 to f^6 . This induces the reachability specification $\psi_{f^5}^1(f^6)$ containing all sequences of rooms in $\mathcal{G}_{f^5}^1$ with final room s_{56} . Now a memoryless strategy for this specification first needs to request a context change from r_{11}^5 to r_{21}^5 . This request, in turn, induces the reachability specification $\psi_{r_{11}^5}^1(r_{21}^5)$ containing all sequences of cells in $\mathcal{G}_{r_{11}^5}$ with final cell q_{43}^5 . A possible first move of the robot to fulfill this specification is from q_{22}^5 to q_{32}^5 . The respective goal states of the two specifications are indicated in dark gray in Fig. 3. \triangleleft

The construction in (21) implies that only a (possibly strict) prefix ξ of a play $\pi \in \phi_\nu^l(\nu')$ needs to be contained in ϕ_ν^l . While this might seem restrictive for non-suffix closed specifications such as safety, one can circumvent this problem by using the idea of “weak until”. Intuitively, one would specify to stay safe, i.e., only visit states from a set Q_{safe} , “until” the context is left. Then (21) checks if the current requested context change can be enforced by staying in safe states. For reachability type specifications, such as the request of the completion of a certain task, this issue does not arise.

Given the above definitions of local specifications, hierarchical reactive games can be constructed from a set of LGGs as follows.

Definition 3: Given a set of local specifications $[\varphi]$ over a set of LGGs $[\mathbb{G}]$ and a set of level 0 initial states $\mathcal{I} \subseteq (X \times Y)$, the tuple $([\mathbb{G}], \mathcal{I}, [\varphi])$ is called a *hierarchical reactive game* (HRG) over $[\mathbb{G}]$. Furthermore, given the set of local initial conditions

$$\mathcal{I}^l(m) := \begin{cases} \{(\alpha_e^{l\uparrow}(x, y), \alpha_s^{l\uparrow}(y)) \mid (x, y) = \mathcal{I}\}, & m = 0 \\ \{\lceil \check{p}_\downarrow^l(m-1) \rceil\}, & m > 0, l < L \\ \text{undefined}, & \text{else}, \end{cases} \quad (22)$$

a set $[\check{p}]_\pi$ is defined to be *winning* (resp. *possibly winning*) for $([\mathbb{G}], \mathcal{I}, [\varphi])$, if for all $l \in [0, L-1]$ holds that

- (i) for all $m \in \text{dom}(\check{\pi}^{l+1})$ (with $m < \text{end}(\check{\pi}^{l+1})$ if $\text{end}(\check{\pi}^{l+1}) < \infty$) there exists a prefix $\xi \sqsubseteq \check{p}_\downarrow^l(m)$ s.t. ξ is winning for $(\mathcal{G}_{\check{y}^{l+1}(m)}^l, \mathcal{I}^l(m), \varphi_{\check{y}^{l+1}(m)}^l)$, and
- (ii) for $m = \text{end}(\check{\pi}^{l+1}) < \infty$ there exists a string $\xi = \check{p}_\downarrow^l(m)$ (resp. $\xi \sqsubseteq \check{p}_\downarrow^l(m)$ or $\check{p}_\downarrow^l(m) \sqsubseteq \xi$) s.t. ξ is winning for $(\mathcal{G}_{\check{y}^{l+1}(m)}^l, \mathcal{I}^l(m), \varphi_{\check{y}^{l+1}(m)}^l)$, and
- (iii) $\check{\pi}^L$ is winning (resp. possibly winning) for $(\mathcal{G}^L, \mathcal{I}^L(0), \varphi^L)$. \triangleleft

V. ASSUME-ADMISSIBLE HIERARCHICAL STRATEGY CONSTRUCTION

Let $([\mathbb{G}], \mathcal{I}, [\varphi])$ be a HRG with initial condition $\mathcal{I} \in (X \times Y)$ and let $[\zeta]$ be a set of local environment assumptions over $[\mathbb{G}]$. Then we want to synthesize a strategy (i.e., a controller) for layer 0 that generates a play whose projection is winning for the set of local system specifications $[\varphi]$ if $[\zeta]$ holds. We assume that $[\varphi]$ and $[\zeta]$ are both ω -regular languages. While in principle one can flatten the game and solve one global game to obtain a solution to this problem, this will be prohibitively expensive. We therefore propose an algorithm that constructs a winning strategy in each local game that is encountered and “stitches together” these winning strategies dynamically. Additionally, one could statically solve and memorize all possibly constructed local games. Our algorithm avoids this expensive construction by only solving games that actually arise online. Hence, our procedure is *dynamic* in that it solves a series of local games in each step starting from the current state — this is conceptually similar to receding horizon control approaches. To incorporate environment assumptions, we use a slightly modified version of the algorithm from Brenguier et al. (2015) to compute an assume-admissible winning strategy for a local game and a local environment assumption. Our procedure treats this algorithm as a black box; in principle, a different strategy synthesis algorithm can be used.

A. Synthesis of Assume-Admissibly Winning Strategies

Assume-admissibly winning strategies for the play $(G, \mathcal{I}, \varphi)$ w.r.t. the assumption ζ can be computed by the algorithm given by Brenguier et al. (2015, Thm. 4) in case φ and ζ are ω -regular objectives. We denote the outcome of this strategy synthesis by $\text{Sol}^{\text{AA}}(G, \mathcal{I}, \varphi, \zeta)$. Whenever the environment does not play admissibly, the definition of assume-admissibly winning strategies does only restrict the behavior of the system to an admissible one. This does not give any guarantees w.r.t. φ in case the environment does not play admissibly. To circumvent this issue we slightly modify the outcome of the available strategy synthesis.

Definition 4: Let $f^{\text{AA}} = \text{Sol}^{\text{AA}}(G, \mathcal{I}, \varphi, \zeta)$ be an assume-admissibly winning strategy, then its associated *possibly winning strategy* f , is defined for all $\pi \in \mathcal{G}$ s.t.

$$f(\pi|_{[0,k]}, x(k+1)) = \begin{cases} f^{\text{AA}}(\pi|_{[0,k]}, x(k+1)), & \pi|_{[0,k]} \cdot (x(k+1), f^{\text{AA}}(\pi|_{[0,k]}, x(k+1))) \in \overline{\varphi} \\ \emptyset, & \text{else.} \end{cases} \quad (23)$$

We define the set of all possibly winning strategies for the game $(G, \mathcal{I}, \varphi)$ w.r.t. ζ by $\text{Sol}(G, \mathcal{I}, \varphi, \zeta)$. \triangleleft

A strategy $f = \text{Sol}(G, \mathcal{I}, \varphi, \zeta)$ blocks whenever the environment forces the play into a state from which the play cannot be won anymore. This implies that all finite plays π compliant with f are possibly winning, i.e. $\pi \in \overline{\varphi}$, even if the environment does not play admissible. However, if it does, the compliant play is winning. This is formalized by the following proposition.

Proposition 3: Given $f = \text{Sol}(G, \mathcal{I}, \varphi, \zeta)$, $g \in \mathcal{S}^e(G)$, it holds for all $\pi \in \mathcal{G}$ that

$$g \in \text{AdmissibleStrategies}(G, \mathcal{I}, \zeta) \Rightarrow f \in \text{WinningStrategies}(G, \mathcal{I}, \varphi, g), \quad (24a)$$

$$\text{and } \left(\begin{array}{l} \pi \in \text{CompliantPlays}(f, \mathcal{I}) \\ \wedge |\pi| < \infty \end{array} \right) \Rightarrow \pi \in \text{WinningPlays}(G, \mathcal{I}, \overline{\varphi}). \quad (24b)$$

Proof: Let $f^{\text{AA}} = \text{Sol}^{\text{AA}}(G, \mathcal{I}, \varphi, \zeta)$ and f its associated possibly winning strategy. Using (4), $g \in \text{AdmissibleStrategies}(G, \mathcal{I}, \zeta)$ implies $f^{\text{AA}} \in \text{WinningStrategies}(G, \mathcal{I}, \varphi, g)$. Using (3), this implies $\pi \in \text{WinningPlays}(G, \mathcal{I}, \overline{\varphi})$. Therefore, the second case in (23) cannot occur and we obtain $f = f^{\text{AA}}$, i.e., $f \in \text{WinningStrategies}(G, \mathcal{I}, \varphi, g)$. Observe that the left side of (24b) implies that the right side of (2) holds for π and f , hence $f(\pi|_{[0, k-1]}, x(k)) \neq \emptyset$ for all $k \in \text{dom}(\pi)$. Using (23), this implies $\pi|_{[0, k]} \in \text{CompliantPlays}(f^{\text{AA}}, \mathcal{I})$ and $\pi|_{[0, k]} \in \overline{\varphi}$, hence, $\pi \in \text{WinningPlays}(G, \mathcal{I}, \overline{\varphi})$. \square

We remark that the algorithm to compute assume-admissible strategies in Brenguier et al. (2015, Thm. 4) can be trivially adapted to ensure Prop. 3, by blocking the game whenever a losing state (one in which there is no winning strategy for the system) is entered.

B. The Strategy Synthesis Algorithm

Recall that we aim to synthesize a strategy (i.e., a controller) for layer 0 that generates a play whose projection is assume-admissible winning for the HRG $([\mathbb{G}], \mathcal{I}, [\varphi])$ w.r.t. $[\zeta]$. Hence, the goal of each computation round of our algorithm is to determine the next system state $y(k+1)$ in layer 0, i.e., to calculate the current control action that needs to be applied to the system. This depends on the environment state $x(k+1)$ in layer 0 which is sensed in the beginning of each such computation round and projected to all layers $l \in [1, L]$ in an “bottom up” fashion. The current state in every layer local game is given by the restriction of $x^l(k+1)$ to the current context and the projection $y^l(k)$ of the last system state. Based on this information, the next step in every layer local game needs to be calculated.

This calculation is challenging due to the interaction between plays in different layers. In particular, a move from system state ν to ν' requested by a strategy in layer $l \in [1, L]$ results in an additional reachability specification for the current local game in layer $l-1$. Furthermore, such an “induced” reachability specification for the local game in layer $l-1$ and context ν might change multiple times, before this context is left. This is due to the fact that an environment state in layer $l > 0$ possibly changes multiple times before a system state change follows, as discussed in the construction of abstract game graphs (see Sec. III-B). Hence, whenever such a specification change occurs, the strategy in layer $l-1$ needs to be re-calculated. The only strategy that is not influenced by this interplay is the highest level strategy, which is computed only once when initializing the algorithm. Once the strategies are updated in a “top down” manner, the controller picks the next move at layer 0 based on the updated strategy for layer 0 and plays it. This changes the states for all higher layers and the algorithm continues with the next computation cycle.

We now describe the algorithm formally.

Algorithm 1 (Strategy Synthesis Procedure): Let $([\mathbb{G}], \mathcal{I}, [\varphi])$ be a HRG with $\mathcal{I} \in (X \times Y)$ and $[\zeta]$ a set of local environment assumptions over $[\mathbb{G}]$. Then the dynamic hierarchical strategy $F = \{f^l\}_{l=0}^L$ for the game $([\mathbb{G}], \mathcal{I}, [\varphi])$ w.r.t. $[\zeta]$ and its compliant play π are iteratively defined as follows:

► Initialization:

▷ Using \mathcal{I}^L as in (22), calculate the assume admissible winning strategy for the highest layer L using

$$h^L = \text{Sol}(G^L, \mathcal{I}^L(0), \varphi^L, \zeta^L). \quad (25a)$$

▷ Initialize the play and the local history, respectively, with

$$\pi = (x(0), y(0)) = \mathcal{I} \quad \text{and} \quad \gamma^l(0) = \pi. \quad (25b)$$

► Iteration for all $k \in \mathbb{N}$:

▷ Sense the environment move

$$x(k+1) \in \delta^0(\pi). \quad (25c)$$

▷ Compute the local environment state $x_{\downarrow}^l(k+1)$ using (6) and (15a), i.e.,

$$x_{\downarrow}^l(k+1) = \mathbf{t}_{y^{l+1}(k)}^l(\alpha_e^{l\uparrow}(x(k+1), y(k))) \quad (25d)$$

for each layer l ;

▷ Iteratively calculate the current strategy by

$$f^L(k) = h^L \quad \text{and} \quad (25e)$$

$$\forall l \in [0, L-1] . f^l(k) = \begin{cases} \emptyset, & \text{GotStuck}^{l+1}(k) \\ h^l(k), & \text{Done}^{l+1}(k) \\ f_{\nu\nu^{l+1}}^l(k), & \text{else} \end{cases} \quad (25f)$$

with

$$\begin{aligned} \nu &:= y^{l+1}(k), \\ \nu^{l+1}(k) &:= f^{l+1}(k)(\check{\gamma}^{l+1}(k), x_{\downarrow}^{l+1}(k+1)), \end{aligned} \quad (25g)$$

$$h^l(k) := \begin{cases} \text{Sol}(G_{\nu}^l, \{\check{\gamma}^l(k)\}, \varphi_{\nu}^l, \zeta_{\nu}^l), & \nu \neq y^{l+1}(k-1) \\ h^l(k-1), & \text{else} \end{cases} \quad (25h)$$

$$f_{\nu\nu^{l+1}}^l(k) = \begin{cases} \text{Sol}(G_{\nu}^l, \{\check{\gamma}^l(k)\}, \phi_{\nu}^l(\nu^{l+1}(k)), \zeta_{\nu}^l), & \left(\begin{array}{l} \nu \neq y^{l+1}(k-1) \\ \vee \nu^{l+1}(k) \neq \nu^{l+1}(k-1) \end{array} \right) \\ f_{\nu\nu^{l+1}}^l(k-1), & \text{else} \end{cases} \quad (25i)$$

and the predicates are defined by

$$\text{Win}^l(k) \Leftrightarrow \check{\gamma}^l(k) \in \begin{cases} \varphi_{\nu}^l, & l \in [0, L-1] \\ \varphi^L, & l = L \end{cases}, \quad (25j)$$

$$\text{Done}^l(k) \Leftrightarrow \left(\begin{array}{l} l = L \vee \text{Done}^{l+1}(k) \\ \wedge \text{Win}^l(k) \\ \wedge (\check{\gamma}^l(k), x_{\downarrow}^l(k+1)) \notin \text{dom}(h^l(k)) \end{array} \right), \quad \text{and} \quad (25k)$$

$$\text{GotStuck}^l(k) \Leftrightarrow \left(\begin{array}{l} \neg \text{Done}^l(k) \\ \wedge (\check{\gamma}^l(k), x_{\downarrow}^l(k+1)) \notin \text{dom}(f^l(k)) \end{array} \right). \quad (25l)$$

▷ Play the next move following the current system strategy for layer $l = 0$

$$y(k+1) = f^0(k)(\check{\gamma}^0(k), x_{\downarrow}^0(k+1)). \quad (25m)$$

▷ Append $(x(k+1), y(k+1))$ to the play giving

$$\pi = (x|_{[0, k+1]}, y|_{[0, k+1]}). \quad (25n)$$

▷ Using (16b), compute the new context restricted history

$$\check{\gamma}^l(k+1) = \lceil \check{p}_{\downarrow}^l \rceil \quad \text{with} \quad \check{p}_{\downarrow}^l \in [\check{p}]_{\pi}. \quad (25o)$$

As discussed before, every computation round k of the construction in (25) starts with the sensing of the next environment move in (25c), giving the full 0-level environment state $x(k+1) = x^0(k+1)$. This state is used to compute the local restricted environment states $x_{\downarrow}^l(k+1)$ for every layer and current context $y^{l+1}(k)$ in (25d). Note that this construction is done “bottom up”.

Thereafter, the selection of the current strategy f^l for every layer and its respective current goal state ν^l are calculated. Observe that this is done “top down”, as ν^l is used to calculate the current reachability specification for the reachability game in layer $l-1$. The construction of f^l in (25f) distinguishes three cases: the play at the highest layer has been won, or the play at the higher layer got stuck, or none of these conditions occurred. We consider these cases separately.

For the first case observe, that the specification of level L might be a set of finite strings and local specifications are sets of finite strings by definition (see Sec. IV-B). Therefore, the play constructed in (25) does not need to be infinite to be winning for $[\varphi]$. If the play in layer L is winning for φ^L and the strategy does not request any other move (denoted by the predicate Done^L in (25k)), then this is communicated downwards using the second line of (25f). In this case all lower level strategies must be winning for local specifications only, using the assume-admissible strategy calculated in (25h).

For the second case, observe that the strategy calculation in (25h) and (25i) does not need to have a solution. Further, even if it has a solution, system strategies are not assumed to be left-total. Hence, there might exist (non-admissible) environment moves that cause a blocking of f without the game being winning. These two situations are modeled by the predicate GotStuck^l in (25k). If such a situation occurs, it is communicated downwards by the first line of (25f) resulting

in $\text{GotStuck}^{l'}$ for all $l' < l$ and therefore an abortion of the game. Intuitively, the first time GotStuck^l occurs, it is because of an “unrealizable” local specification. We introduce a fourth predicate

$$\text{UnRealizable}^l(k) \Leftrightarrow \begin{cases} \text{GotStuck}^l(k), & l = L \\ \neg \text{GotStuck}^{l+1}(k) \wedge \text{GotStuck}^l(k), & l < L \end{cases} \quad (26)$$

to remember the first layer at which the controller got stuck. We will show in Sec. V-C that an unrealizable specification is the only reason for a non-winning play constructed in (25) to be aborted.

In the third case, i.e., if neither GotStuck^l nor Done^{l+1} is true, the strategy for level l is calculated by (25i) using again two subcases. In the first subcase, either a new context was entered (resulting in a new local game) or the “top down induced” reachability specification has changed (due to a change of ν^l caused by a new environment state in layer $l+1$). In this case the strategy for level l needs to be re-calculated. However, if neither of these two situations occurs, the strategy from the previous time step can be used, avoiding unnecessary re-computations.

After the strategy construction in (25f)-(25l), the system state is updated to $y(k+1)$, using the currently selected lowest level strategy $f^0(k)$ in (25m). Hence, (25f)-(25l) only utilize the hierarchical structure of the game graph to compute $f^0(k)$, which is the only control action that is actually applied to the system, e.g., the robot in our example. Then $(x(k+1), y(k+1))$ is appended to the constructed play π . As intuitively assumed, such plays π generated by Alg. 1 up to length k are plays in G , i.e., $\pi \in \mathcal{G}$, as shown in the following proposition. Observe, that this implies that also $\tilde{\pi}^l \in \mathcal{G}^l$ for all $l \in [0, L]$ (from Prop. 1) and $\tilde{p}_\downarrow^l(m) \in \mathcal{G}_{\tilde{y}^{l+1}(m)}^l$ for all $l \in [0, L-1]$ and $m \in \text{dom}(\tilde{\pi}^{l+1})$ (from Prop. 2).

Proposition 4: Let π be a play computed in Alg. 1. Then $\pi \in \mathcal{G}$.

Proof: It follows from (25c) and (25m) that

$$\forall k \in \text{dom}^+(\pi) . \left(\begin{array}{l} x(k) \in \delta(x(k-1), y(k-1)) \\ \wedge y(k) = f^0(k-1)(\tilde{\gamma}^0(k-1), x_\downarrow^0(k)) \end{array} \right), \quad (27)$$

implying $f^0(k-1) \neq \emptyset$ for all $k \in \text{dom}^+(\pi)$. Therefore, (25f)-(25l) imply that $f^0(k-1)$ is a system strategy over $\mathcal{G}_{y^1(k-1)}^0$ and the definition of the latter in Sec. II gives $f^0(k-1)(\tilde{\gamma}^0(k-1), x_\downarrow^0(k)) \in \rho_{y^1(k-1)}^0(x_\downarrow^0(k-1), [\tilde{\gamma}^0(k-1)]_2)$. Now observe from (25o), (16b) and (8) that $[\tilde{\gamma}^0(k-1)]_2 = y^0(k-1)$. Now using $\rho_{y^1(k-1)}^0 \subseteq \rho^0$ from Ass. 1 along with this observation, we see that (27) actually implies (1), hence $\pi \in \mathcal{G}$. \square

We call a play π calculated in (25) up to length $k = |\pi|$ *maximal* if

$$k < \infty \Rightarrow (\tilde{\gamma}^0(k), x_\downarrow^0(k+1)) \notin \text{dom}(f^0(k)). \quad (28)$$

One round of the construction in (25) is ended by calculating the current local histories $\tilde{\gamma}^l(k+1)$ for every layer. Intuitively, $\tilde{\gamma}^l(k+1)$ models the part of $\tilde{\pi}^l$ generated after the last context change in layer l and is therefore equivalent to $[\tilde{p}_\downarrow^l]$. These histories are used in the calculation of assume-admissible strategies to ensure that a re-computation of a strategy within one context does result in a continuation of the already generated string w.r.t. the given specification.

While the local system strategies $f^l(k)$ are explicitly calculated for every time step k in (25f)-(25l), the local environment strategies $g^l(k)$ are only given implicitly by the observed environment move (25c) and its abstraction to every layer l . Formally, a play π calculated in (25) was played against an admissible environment strategy if for all $l \in [0, L-1]$, $m \in \text{dom}(\tilde{\pi}^l)$ exists an environment strategy $g_{\tilde{y}^{l+1}(m)}^l \in \text{AdmissibleStrategies}(G_{\tilde{y}^{l+1}(m)}^l, \mathcal{I}^l(m), \zeta_{\tilde{y}^{l+1}(m)}^l)$ s.t. $\tilde{p}_\downarrow^l(m) \in \text{CompliantPlays}(G_{\tilde{y}^{l+1}(m)}^l, g_{\tilde{y}^{l+1}(m)}^l)$ and for layer L exists $g^L \in \text{AdmissibleStrategies}(G^L, \mathcal{I}^L(0), \zeta^L)$ s.t. $\tilde{\pi}^L \in \text{CompliantPlays}(G^L, g^L)$. If this holds, we call π an *environment admissible play*.

Example 9: Consider the play π whose y -component is depicted by filled cycles in Fig. 1 (bottom) and (for simplicity) the static environment used in Expl. 7, where we use $o = \{q_{24}^5, q_{25}^5, q_{63}^5\}$ and $o_\downarrow = \{q_{24}^5, q_{25}^5\}$ for notational convenience. In this game the only objective is to reach q_{63}^6 in r_{21}^6 and f^6 . This implies that $[\varphi]$ contains only empty sets except for

$$\varphi^2 = \{\perp\} \times \{f^5 f^6\}, \quad \varphi_{f^6}^1 = \{\perp\} \times R^* \cdot \{r_{21}^6\}, \quad \text{and} \quad \varphi_{r_{21}^6}^0 = \{\perp\} \times Q^* \cdot \{q_{21}^6\}.$$

To illustrate Alg. 1 we pick $k = 2$, i.e., π was generated for 3 time steps and we are now calculating $\pi(3) = (x(3), y(3))$ using (25).

First recall from Expl. 7 that

$$\begin{aligned} \pi(2) &= (o, q_{33}^5), \quad \pi^1(2) = \pi^1(0) = (\{d\}, r_{11}^5), \quad \pi^2(2) = \pi^2(0) = (\{\perp\}, f^5), \quad \text{and} \\ \tilde{\gamma}^0(2) &= (o_\downarrow, q_{22}^5)(o_\downarrow, q_{23}^5)(o_\downarrow, q_{33}^5), \quad \tilde{\gamma}^1(2) = (\{d\}, r_{11}^5), \quad \tilde{\gamma}^2(2) = (\{\perp\}, f^5). \end{aligned}$$

We furthermore assume that the strategy calculation for $k = 0$ resulted in the requested moves depicted by the arrows in Fig. 3 (middle and top). With this initialization we obtain the following steps of the algorithm.

- ▷ Due to the static environment assumption, (25c) gives $x(k+1) = x(3) = o$.
- ▷ Applying (25d) yields $x_\downarrow^0(3) = o_\downarrow$, $x_\downarrow^1(3) = \{d\}$ and $x_\downarrow^2(3) = \{\perp\}$.
- ▷ First, (25e) and (25e) imply $f^2(2) \neq \emptyset$, $\nu^2(2) = \nu^2(1) = f^6$ and $\neg \text{Done}^2(2)$. Therefore, (25i) and (25f) imply $f^1(2) =$

$f_{f^5 f^6}^1(0) \neq \emptyset$, $\nu^1(2) = \nu^1(1) = r_{11}^5$ and $\neg \text{Done}^1(2)$. With this, the lowest level strategy is given by $f^0(2) = f_{r_{11}^5, r_{21}^5}^0(0)$.
 \triangleright As we assume a static environment and no obstacles block the way between the robot and the exit to room r_{21}^5 , we assume that $f_{r_{11}^5, r_{21}^5}^0$ is a shortest path strategy and (25m) gives $y(k+1) = y(3) = q_{43}^5$.
 \triangleright Observe, that a context change has occurred during this step, i.e., (25o) gives

$$\check{\gamma}^0(3) = (x_{\downarrow}^0(3), y(3)) = (o_{\downarrow}, q_{43}^5), \check{\gamma}^1(3) = (\{d\}, r_{11}^5)(\{d\}, r_{21}^5), \check{\gamma}^2(3) = (\{\perp\}, f^5).$$

With this local history the next iteration of the algorithm is started. For the assumed very simple static environment, Alg. 1 will never get stuck. Observe, that once we reach floor f^6 , the level 2 game is won and Done^2 is true. In this case h^1 will be calculated w.r.t. the specification $\varphi_{f^6}^1$. If in addition r_{21}^6 is reached, Done^1 is also set to true and h^0 is calculated. After one more time step also Done^0 is true and the algorithm terminates. The generated play is obviously winning for $[\varphi]$. \triangleleft

C. Soundness

In this section we prove three different soundness results for the play constructed in Alg. 1. Intuitively, Alg. 1 is sound if a play π calculated in (25) is winning for the HRG $([\mathbb{G}], \mathcal{I}, [\varphi])$ if all generated local specifications are realizable and the environment plays admissible w.r.t. $[\zeta]$, which will be proven last in Thm. 3. As a first intermediate result we show that the only two reasons for a maximal play to terminate are actually that (i) a current local specification is not realizable or (ii) the play is already winning given a finite winning condition in layer L .

Theorem 1: Let π be a maximal play computed by (25). Then it holds that

$$|\pi| < \infty \Leftrightarrow \left(\forall l \in [0, L] . \text{Done}^l(\text{end}(\pi)) \right) \vee \left(\exists l \in [0, L] . \text{UnRealizable}^l(\text{end}(\pi)) \right). \quad (29)$$

Proof: To prove this theorem we need that

$$\left(\exists l \in [0, L] . \text{UnRealizable}^l(\text{end}(\pi)) \right) \Leftrightarrow \text{GotStuck}^0(\text{end}(\pi)) \quad (30a)$$

which is proven for all $k \in \text{dom}(\pi)$ in Lem. 5 (see App.). Furthermore, as we assume environment strategies to be left-total, (25c) can always be computed. Hence, π becomes finite while being maximal iff (25m) cannot be evaluated, i.e.,

$$\text{end}(\pi) < \infty \Leftrightarrow (\check{\gamma}^0(\text{end}(\pi)), x_{\downarrow}^0(\text{end}(\pi) + 1)) \notin \text{dom}(f^0(\text{end}(\pi))). \quad (30b)$$

Now we pick $k = \text{end}(\pi)$ and prove both directions separately.

“ \Rightarrow ” Using (30b) and (25l) implies that either (i) $\neg \text{Done}^0(k)$ and $\text{GotStuck}^0(k)$, or (ii) $\text{Done}^0(k)$. Using (30a), (i) implies $\langle (29).\text{right}.2 \rangle$. As $\text{Done}^0(k)$ implies $\forall l \in [0, L] . \text{Done}^l(k)$ (from (25k)), (ii) implies $\langle (29).\text{right}.1 \rangle$.

“ \Leftarrow ” If $\langle (29).\text{right}.2 \rangle$ is true, it follows from (30a) that $\text{GotStuck}^0(k)$ and $\neg \text{Done}^0(k)$ (see the proof of Lem. 5). Hence, (29) and (30b) implies $\langle (29).\text{left} \rangle$. If $\langle (29).\text{right}.1 \rangle$ is true, we know from (25f) that $f^0(k) = h^0(k)$. Therefore, $\langle (25k).\text{right}.3 \rangle$ and (30b) implies $\langle (29).\text{left} \rangle$. \square

While the second case in Thm. 1 is not desired w.r.t. the goal of constructing a winning play, it can usually not be avoided in a realistic scenario as we can (i) not enforce the environment to play admissible and (ii) checking feasibility of all possibly occurring local games before startup might not be appropriate, as this set might be very large. However, Alg. 1 ensures that if this situation occurs, the local specifications are not falsified up to this point. This is formalized by the notion of possibly winning, which ensures that generated finite plays always stay in the prefix closure of the considered local specifications.

Theorem 2: Given the preliminaries of Alg. 1, let π be the play computed by (25) up to length k , and $[\check{p}]_{\pi}$ its set of local projected play sequences. Then $[\check{p}]_{\pi}$ is possibly winning for $([\mathbb{G}], \mathcal{I}, [\varphi])$.

Proof: We have two important observations that we use in this proof. First, it holds for all $l \in [0, L]$ and $m \in \text{dom}^+(\check{\pi}^l)$ that

$$\left(\begin{array}{l} \check{x}_{\downarrow}^l(m) \in \delta_{\check{y}^l(m-1)}^l(\check{x}_{\downarrow}^l(m-1), \check{y}^l(m-1)) \\ \wedge \check{y}^l(m) = f^l(\kappa^l(m) - 1)(\check{\gamma}^l(\kappa^l(m) - 1), \check{x}_{\downarrow}^l(m)) \end{array} \right) \quad (31a)$$

as proven in Lem. 8 (see App.). Second, it holds for all $l \in [0, L-1]$ and $m \in \text{dom}^+(\check{\pi}^{l+1})$ that

$$\check{p}_{\downarrow}^l(m-1) \in \phi_{\check{y}^{l+1}(m-1)}^l(\check{y}^{l+1}(m)) \quad (31b)$$

and for $m = \text{end}(\check{\pi}^{l+1})$ there exists $\nu' \in \text{Post}^{l+1}(\check{y}^{l+1}(m))$ s.t.

$$\check{p}_{\downarrow}^l(m) \in \overline{\phi_{\check{y}^{l+1}(m)}^l(\nu')}, \quad (31c)$$

⁴To simplify notation we denote by $\langle (\#).\text{right}.n \rangle$ (resp. $\langle (\#).\text{left}.n \rangle$) the n th statement on the right (resp. left) side of the implication/equivalence relation in equation (#).

as proven in Lem. 9 (see App.).

Recall from Prop. 4 that $\pi \in \mathcal{G}$, hence Prop. 2 implies $\check{p}_\downarrow^l(m) \in \mathcal{G}_{\check{y}^{l+1}(m)}^l$ and (16a) obviously gives $\check{p}_\downarrow^l(m)|_{[0,0]} = [\check{p}_\downarrow^l(m-1)] = \mathcal{I}^l(m)$ for all $m \in \text{dom}^+(\check{\pi}^{l+1})$. As (31b) holds, (21) implies

$$\exists \xi \in \overline{\{\check{p}_\downarrow^l(m-1)\}} \cdot \xi \in \varphi_{\check{y}^{l+1}(m-1)}^l. \quad (32a)$$

Now consider $m = \text{end}(\check{\pi}^{l+1})$. As (31c) holds, (21) implies that either

$$\check{p}_\downarrow^l(m) \in \overline{\varphi_{\check{y}^{l+1}(m)}^l} \quad \text{or} \quad \exists \xi \in \overline{\{\check{p}_\downarrow^l(m)\}} \cdot \xi \in \varphi_{\check{y}^{l+1}(m)}^l. \quad (32b)$$

Using the definitions of winning from Sec. II, (32a)-(32b) imply that conditions (i)-(ii) for possibly winning HRGs from Sec. IV-B hold. To prove condition (iii), observe from (25e) that $\forall k \in \mathbb{N} \cdot f^L(k) = h^L$. Furthermore, recall from the definition of $[\check{p}]_\pi$ that $\check{p}_\downarrow^L(0) = \check{\pi}^L$ and $\text{end}(\check{p}_\downarrow^L) = 0$ and therefore $\check{\gamma}^L(\kappa^L(m) - 1) = \check{\pi}^L|_{[0, \kappa^L(m) - 1]}$. Using these observations in (31a), it follows that (2) holds for $\check{\pi}^L$ w.r.t. h^L and $\mathcal{I}^L(0)$, implying $\pi^L \in \overline{\text{CompliantPlays}(h^L, \mathcal{I}^L(0))}$. As $h^L = \text{Sol}(G^L, \mathcal{I}^L(0), \varphi^L, \zeta^L)$ and $\check{\pi}^L \in \mathcal{G}^L$ (from Prop. 4 and Prop. 1), it follows from (24b) in Prop. 3 that $\check{\pi}^L$ is possibly winning for $(\mathcal{G}^L, \mathcal{I}^L(0), \varphi^L)$. \square

We now prove the main result of this paper, namely that maximal plays π calculated by Alg. 1 (finite and infinite) are actually winning for $([\mathbb{G}], \mathcal{I}, [\varphi])$ if the environment plays admissible and all constructed local plays have a solution, i.e.,

$$\forall k \in \text{dom}(\pi), l \in [0, L] \cdot \neg \text{UnRealizable}^l(k). \quad (33)$$

Theorem 3: Let π be a maximal and environment admissible play computed by (25) s.t. (33) holds and let $[\check{p}]_\pi$ be its set of local play sequences. Then $[\check{p}]_\pi$ is winning for $([\mathbb{G}], \mathcal{I}, [\varphi])$.

Proof: In this proof we use the following two observations

$$\left(\forall k \in \text{dom}(\pi), l \in [0, L] \cdot \neg \text{Done}^l(k) \right) \Leftrightarrow (|\pi| = \infty) \Leftrightarrow (\forall l \in [0, L] \cdot |\check{\pi}^l| = \infty), \quad (34a)$$

$$\left(\forall l \in [0, L] \cdot \text{Done}^l(\text{end}(\pi)) \right) \Leftrightarrow (|\pi| < \infty) \Leftrightarrow (\forall l \in [0, L] \cdot |\check{\pi}^l| < \infty). \quad (34b)$$

where (34a) was proven in Lem. 11 (see App.), the left side of (34b) follows from Thm. 1 and (33), and the right side of (34b) is a simple consequence from the definition of projections in (8). Hence, we generally have two cases to consider when proving the three conditions for winning HRGs from Sec. IV-B.

First observe that condition (i) is equivalent for winning and possibly winning, no matter whether π is finite or not. It therefore follows directly from Thm. 2. Furthermore, condition (ii) only needs to be proven if $|\check{\pi}^{l+1}| < \infty$ and recall that for this case Thm. 2 shows that $\check{p}_\downarrow^l(\text{end}(\check{\pi}^{l+1}))$ is possibly winning for $(\mathcal{G}_{\check{y}^{l+1}(m)}^l, [\check{p}_\downarrow^l(m-1)], \varphi_{\check{y}^{l+1}(m)}^l)$ for all $l \in [0, L]$. Now observe from (34b) that $\text{Done}^l(\text{end}(\pi))$ which implies from (25k) and (25j) that $\check{p}_\downarrow^l(\text{end}(\check{\pi}^{l+1})) = \check{\gamma}^l(\text{end}(\pi)) \in \varphi_{\check{y}^{l+1}(m)}^l$, where the first equality follows from (25o) and (16). This obviously implies that $\check{p}_\downarrow^l(\text{end}(\check{\pi}^{l+1}))$ is winning in the above game. For finite plays, this reasoning also proves condition (iii). We therefore assume $|\check{\pi}^L| = \infty$ and recall from the proof of Thm. 2 that (2) holds for $\check{\pi}^L$ w.r.t. h^L and $\mathcal{I}^L(0)$. As $|\check{\pi}^L| = \infty$ we have $\check{\pi}^L \in \text{CompliantPlays}(h^L, \mathcal{I}^L(0))$. As $h^L = \text{Sol}(G^L, \mathcal{I}^L(0), \varphi^L, \zeta^L)$ and $\check{\pi}^L \in \mathcal{G}^L$ (from Prop. 4 and Prop. 1) and $g^L \in \text{AdmissibleStrategies}(G^L, \mathcal{I}^L(0), \varphi^L, \zeta^L)$, it follows from (24b) in Prop. 3 that $\check{\pi}^L$ is winning for $(\mathcal{G}^L, \mathcal{I}^L(0), \varphi^L)$. \square

The important difference between Thm. 2 and Thm. 3 is that environment admissible infinite plays can only be generated if layer L does not win in finite time, i.e., $\neg \text{Done}^L(k)$ for all $k \in \text{dom}(\check{\pi}^L)$. If the environment does not play admissible, infinite plays can also be generated if $\text{Done}^L(k)$ is true, as the environment might never “help” to reach the specification (i.e., does not play admissible) but also never moves to a losing state (i.e., causing the game to be aborted).

Remark 1: It should be noted that the algorithm in Alg. 1 works identically if we use a “usual” synthesis techniques to calculate winning (instead of assume-admissibly winning) strategies in $\text{Sol}(\cdot)$ (i.e., a procedure to solve the unconstrained synthesis problem). Such a procedure is obtained, e.g., from the methods by Zielonka (1998); Emerson and Jutla (1991) for general ω -regular conditions, or more specialized procedures for co-safe properties (given by sets of finite-length plays) by Kupferman and Vardi (2001); Ehlers and Finkbeiner (2011); Kupferman and Weiner (2012). This outlines the modularity of our approach w.r.t. the actual strategy synthesis routine used in local games. However, it should be noted that in realistic scenarios, local games will usually not have winning strategies against a purely adversarial environment. Nevertheless, if the game gets stuck due to such an unrealizable sub-game, the result from Thm. 2 still holds, i.e., the specification is not violated in this case.

D. Comments on Completeness

Intuitively, the synthesis procedure given in Alg. 1 is complete if, whenever there exists a strategy \hat{f} over the game graph G s.t. all plays $\hat{\pi} \in \mathcal{G}$ compliant with \hat{f} induce a set of local play sequences that are winning for $([\mathbb{G}], \mathcal{I}, [\varphi])$ (if the environment plays an admissible strategy), then there exists a hierarchical strategy F s.t. its compliant play π generated by (25) induces projected plays that are also winning for $([\mathbb{G}], \mathcal{I}, [\varphi])$ (if the environment plays an admissible strategy).

Unfortunately, this statement is not true. The major problem arises from the fact that assume-admissibly winning strategies are usually not unique for a particular game. Therefore, using one particular strategy calculated by $\text{Sol}(\cdot)$ disregards other winning plays. This has two important consequences. First, a move of the current layer l strategy cannot be revised if the current layer $l - 1$ game is not realizable for the corresponding reachability specification, even if there exists a different possibly winning extension in layer l . In our robot example, this corresponds to the case where the robot is in a particular room r with two adjacent rooms r' and r'' , where visiting either of them is winning. Now the current strategy for the room layer deterministically picks room r' . If the way towards room r' is blocked by a static obstacle, the game in layer 0 and context r does not have a solution and the play gets stuck.

This problem also arises in reverse layer interaction, as assume-admissibly winning strategies are only ensured to be winning against a “local” admissible environment strategy. They do not consider admissible environment moves in higher layers that might cause specification changes in the current layer. Hence, the local strategy synthesis might pick a strategy that leads the play to a region of the state space which is losing for a different specification that might occur later in this game due to such an admissible environment move in a higher layer. In the above example this would correspond to the case that the door to room r' gets closed which is visible to layer 1 and therefore causes the strategy to request the robot to move to room r'' , instead. Now assume that the way towards both r' and r'' was unblocked initially. Given the specification to reach r' the robot might pick one of two passages which allow to reach r' but the selected one is too narrow for the robot to turn. When the specification changes, the robot cannot turn and approach r'' , hence the game in layer 0 and context r does not have a solution and the play gets stuck. Taking these interactions into account when synthesizing local assume-admissible winning strategies is a promising idea for future work to obtain a complete algorithm. This would also reduce blocking situations which are caused by this interplay.

Completeness holds in the special case of a trivial environment (which has no choice of moves) and the strategy only picks one among the choice of system moves (as e.g. in Kloetzer and Belta, 2008; Vasile and Belta, 2014). However, in this case, one can compute a strategy statically using a dynamic programming procedure similar to context free reachability (see Reps et al., 1995; Alur et al., 2003).

VI. CONCLUSION

We have shown in this paper how a large-scale reactive controller synthesis problem with intrinsic *hierarchy* and *locality* can be modeled as a hierarchical two player game over a set of local game graphs w.r.t. to a set of local strategies on multiple, interacting abstraction layers. We have proposed a reactive controller synthesis algorithm for such hierarchical games that allows for *dynamic specification changes* at each step of the play which is recalculated online in every step. This re-calculation becomes computationally tractable by the proposed decomposition. We have shown that our algorithm is sound: whenever the environment meets its assumptions and all dynamically generated local games have a solution, the controller synthesis algorithm generates a winning hierarchical play for a given specification. If these assumptions do not hold, the algorithm terminates but the generated finite play does not violate the specification up to this point.

APPENDIX

Lemma 1: Let π be a play and κ^l its timescale transformation for level $l \in [0, L]$. For all $l \in [0, L - 1]$, we have

$$\forall k \in \text{dom}(\kappa^{l+1}) . \exists m \in \text{dom}(\kappa^l) . \kappa^{l+1}(k) = \kappa^l(m) \quad \text{and} \quad \lceil \kappa^l \rceil \geq \lceil \kappa^{l+1} \rceil.$$

Proof: We prove both statements by contradiction.

Take $k \in \text{dom}(\kappa^{l+1})$ and define $n = \kappa^{l+1}(k)$. Assume that there exists no $m \in \text{dom}(\kappa^l)$ s.t. $n = \kappa^l(m)$. This implies, by the definition of κ^l in (7), that $y^l(n-1) = y^l(n)$. However, this implies (by definition of layers) that $y^{l+1}(n-1) = y^{l+1}(n)$, which is a contradiction as the assumption $n = \kappa^{l+1}(k)$ implies (from (7)) that $y^{l+1}(n-1) \neq y^{l+1}(n)$. Assume that there exists a $k \in \text{dom}(\kappa^{l+1})$ s.t. $k > \lceil \kappa^l \rceil$ and $n = \kappa^{l+1}(k)$. As before, this implies $y^l(n-1) = y^l(n)$ and hence $y^{l+1}(n-1) = y^{l+1}(n)$ which is a contradiction to the assumption that $k \in \text{dom}(\kappa^{l+1})$. \square

Lemma 2: For each game G , each play π of G and each $l \in [0, L]$, we have

$$\forall m \in \text{dom}(\check{\pi}^l), n \in (\kappa^l(m), \kappa^l(m+1)] . \left(\begin{array}{l} x^l(n) \in \delta^l(\check{x}^l(m), \check{y}^l(m)) \\ \wedge \check{y}^l(m+1) \in \rho^l(\check{x}^l(m+1), \check{y}^l(m)) \end{array} \right). \quad (35)$$

Proof: Pick $l \in [1, L]$ and $m \in \text{dom}(\check{\pi}^l)$ s.t. $m < \text{end}(\kappa^l)$ and $\pi' = \pi|_{[\kappa^l(m), \text{end}(\pi)]}$ and $\pi'' = \pi|_{[\kappa^l(m+1)-1, \text{end}(\pi)]}$. Observe that $\pi', \pi'' \in \mathcal{G}$ by definition and we denote by κ'^l and κ''^l their respective timescale transformations defined via (7). Observe that $m < \text{end}(\kappa^l)$ implies $\text{end}(\kappa'^l), \text{end}(\kappa''^l) > 0$. We therefore obviously have $n \in (0, \kappa'^l(1)]$ and observe from the construction of π' that

$$\begin{aligned} \pi'^l(\kappa'^l(0)) &= \pi^l(\kappa^l(m)) = (\check{x}^l(m), \check{y}^l(m)) \quad \text{and} \\ \pi'^l(\kappa'^l(1)) &= \pi^l(\kappa^l(m+1)) = (\check{x}^l(m+1), \check{y}^l(m+1)). \end{aligned}$$

With this it immediately follows from (10a) that $x^l(n) \in \delta^l(\check{x}^l(m), \check{y}^l(m))$. Observe that $m < \text{end}(\kappa^l)$ implies that $\check{y}^l(m) \neq \check{y}^l(m+1)$. It furthermore follows from (7) that $y^l(\kappa^l(m+1)-1) = \check{y}^l(m)$. Using these observations we have

$$\begin{aligned} \pi^m(\kappa^m(1)-1) &= \pi^l(\kappa^l(m+1)-1) = (x^l(\kappa^m(1)-1), \check{y}^l(m)) \\ \text{and } \pi^m(\kappa^m(1)) &= \pi^l(\kappa^l(m+1)) = (\check{x}^l(m+1), \check{y}^l(m+1)). \end{aligned}$$

With this it immediately follows from (10b) that $\check{y}^l(m+1) \in \rho^l(\check{x}^l(m+1), \check{y}^l(m))$. \square

Lemma 3: Let $[\mathbb{G}]$ be a set of LGGs and \mathcal{G}_y^l the set of plays in G_y^l . Furthermore, let $\pi \in \mathcal{G}$ and $[\check{p}]_\pi$ its induced set of local projected play sequences. Then it holds for all $l \in [0, L-1]$ and $m \in \text{dom}^+(\check{\pi}^{l+1})$ that

$$\left(\begin{array}{l} \forall k \in [\kappa_l^{l+1}(m-1), \kappa_l^{l+1}(m)) \cdot \check{y}^l(k) \in Y_{\check{y}^{l+1}(m-1)}^l \\ \wedge \check{y}^l(\kappa_l^{l+1}(m)) \in Y_{\check{y}^{l+1}(m-1)}^l \cap Y_{\check{y}^{l+1}(m)}^l \\ \wedge \check{p}_\downarrow^l(m-1) \in \mathcal{G}_{\check{y}^{l+1}(m-1)}^l \end{array} \right) \quad (36a)$$

and for all $l \in [0, L-1]$ that

$$\left(\begin{array}{l} \forall k \in [\kappa_l^{l+1}(\text{end}(\check{\pi}^{l+1})), \text{end}(\check{\pi}^l)] \cdot \check{y}^l(k) \in Y_{[\check{y}^{l+1}]}^l \\ \wedge [\check{p}_\downarrow^l] \in \mathcal{G}_{[\check{y}^{l+1}]}^l \end{array} \right). \quad (36b)$$

Proof: As the proof of (36b) is a simplified version of the proof for (36a), we only give the latter. We fix $l \in [0, L-1]$, $m \in \text{dom}(\kappa^{l+1})$ and $k \in [\kappa_l^{l+1}(m-1), \kappa_l^{l+1}(m))$ and prove all lines of the statement separately. To simplify notation we use $\nu := \check{y}^{l+1}(m-1)$ and $\nu' := \check{y}^{l+1}(m)$.

► Pick $r := \kappa^l(k)$ and $r' := \kappa^{l+1}(m)$ and observe that $r \in [\kappa^{l+1}(m-1), \kappa^{l+1}(m))$. With this choice, (7), (8) and (5) imply

$$y^{l+1}(r) = \nu \neq \nu' = y^{l+1}(r'), \quad (37a)$$

$$y^{l+1}(r) = \alpha_s^{l+1}(y^l(r)) \text{ and } y^{l+1}(r') = \alpha_s^{l+1}(y^l(r')). \quad (37b)$$

Substituting $y^l(r) = \check{y}^l(k)$ and $y^l(r') = \check{y}^l(\kappa_l^{l+1}(m))$ in (37b) and using (12c) gives

$$\check{y}^l(k) \in Y_\nu^l, \quad \text{and} \quad \check{y}^l(\kappa_l^{l+1}(m)) \in Y_{\nu'}^l, \quad (38)$$

where the left side of (38) proves the first line of (36a).

► Recall from Prop. 4 that $\check{\pi}^l \in \mathcal{G}^l$. Using Def. 1 this implies that

$$\check{x}^l(k+1) \in \delta^l(\check{x}^l(k), \check{y}^l(k)) \quad \text{and} \quad \check{y}^l(k+1) \in \rho^l(\check{x}^l(k+1), \check{y}^l(k)). \quad (39a)$$

Using the left side of (38) and Ass. 1, (39a) implies

$$\check{y}^l(k+1) \in \rho^l(\tau_\nu^l(\check{x}^l(k+1)), \check{y}^l(k)) = \rho^l(\check{x}_\downarrow^l(k+1), \check{y}^l(k)). \quad (39b)$$

As $\check{x}_\downarrow^l(k+1) = \tau_\nu^l(\check{x}^l(k+1)) \in X_\nu^l$ (from (12a)) it follows from (39b) and (12d) that

$$\check{y}^l(\kappa_l^{l+1}(m)) \in Y_\nu^l. \quad (39c)$$

Combining (39c) with the right side of (38) proves the second line of (36a).

► Using (38), (39c), (12a) and (39a) in (13) implies that

$$\check{x}_\downarrow^l(k+1) \in \delta_\nu^l(\check{x}_\downarrow^l(k), \check{y}^l(k)) \quad \text{and} \quad \check{y}^l(k+1) \in \rho_\nu^l(\check{x}_\downarrow^l(k+1), \check{y}^l(k)), \quad (40)$$

hence, the third line of (36a) holds. \square

Lemma 4: Let π be a maximal play computed by (25). Then it holds for all $l \in [0, L-1]$ and $k \in \text{dom}(\pi)$ that

$$\left(\neg \text{UnRealizable}^l(k) \right) \Leftrightarrow (\check{\gamma}^l(k), x_\downarrow^l(k+1)) \in \text{dom}(f^l(k)) \quad (41)$$

if $\neg \text{Done}^{l+1}(k)$.

Proof: “ \Rightarrow ”: The left side of (41) and (26) implies $\neg \text{GotStuck}^l(k)$ and $\neg \text{Done}^{l+1}(k)$ implies $\neg \text{Done}^l(k)$ from (25k). Using both observations in (25l) implies $(\check{\gamma}^l(k), x_\downarrow^l(k+1)) \in \text{dom}(f^l(k))$.

“ \Leftarrow ”: The right side of (41) implies $f^l(k) \neq \emptyset$. Therefore, it follows from (25f) that $\neg \text{GotStuck}^{l+1}(k)$ and (as $\neg \text{Done}^l(k)$) from (25l) $\neg \text{GotStuck}^l(k)$. Using both observations in (26) also gives $\neg \text{UnRealizable}^l(k)$. \square

Lemma 5: Let π be a maximal play computed by (25). Then it holds for all $k \in \text{dom}(\pi)$ that

$$(\exists l \in [0, L] \cdot \text{UnRealizable}^l(k)) \Leftrightarrow \text{GotStuck}^0(k). \quad (42)$$

Proof: “ \Rightarrow ”: Pick l s.t. $\text{UnRealizable}^l(k)$ and observe that this implies $\text{GotStuck}^l(k)$ (from (26)) and hence $\neg \text{Done}^l(k)$ (from (25l)). Using the first line of (25f) this implies $f^{l-1}(k) = \emptyset$. As $\neg \text{Done}^l(k)$ also implies $\neg \text{Done}^{l-1}(k)$ from (25k) it

follows from (25l) that $\text{GotStuck}^{l-1}(k)$ is true (i.e., $\text{GotStuck}^l(k) \Rightarrow \text{GotStuck}^{l-1}(k)$). Applying this reasoning repetitively we eventually obtain $\text{GotStuck}^0(k)$.

“ \Leftarrow ”: Using (26), $\text{GotStuck}^0(k)$ implies that the right side of (41) in Lem. 4 is false. Hence, either $\text{UnRealizable}^0(k)$ or $\text{GotStuck}^1(k)$ is true. If UnRealizable^0 is true the statement is proven. We therefore assume that $\text{GotStuck}^1(k)$ is true. We can reuse the same reasoning to either eventually get UnRealizable^l for some $l \in [0, L]$ (what proves the statement) or reach $\text{GotStuck}^L(k)$. However, it follows from (26) that the latter is equivalent to UnRealizable^L , what proves the statement. \square

Lemma 6: Let $\pi = (x, y) \in \mathcal{G}_\nu^l$ for some $\nu \in Y^{l+1}$ s.t. $y(0) \in Y_{\nu_l}^l$ and $\psi_\nu^l(\nu')$ with $\nu' \in Y^{l+1}$, $\nu \neq \nu'$ as in (20). Then it holds that

$$\pi \in \psi_\nu^l(\nu') \Leftrightarrow \begin{pmatrix} \text{end}(\pi) < \infty \\ \wedge \forall k < \text{end}(\pi) . y(k) \in Y_{\nu_l}^l \\ \wedge [y] \in Y_{\nu_l}^l \cap Y_{\nu'}^l \end{pmatrix}. \quad (43)$$

Proof: “ \Leftarrow ” $\langle (43).\text{right}.1 \rangle$ and $\langle (43).\text{right}.3 \rangle$ immediately imply that $\pi \in \psi_\nu^l(\nu')$ (from the first line of (20)). “ \Rightarrow ” $\langle (43).\text{right}.2 \rangle$ is the only non-obvious conclusion from (20). Recall that $\pi \in \mathcal{G}_\nu^l$ and $y(0) \in Y_{\nu_l}^l$. Therefore it holds for all $r \leq \text{end}(\pi)$ that $y_\nu^l(r) \in Y_{\nu_l}^l \cup Y_{\nu'}^l$. Now assume that there exists $r' < \text{end}(\pi)$ s.t. $y(r') \in Y_{\nu_l}^l$. Using (12b) this would imply $y(r') \notin Y_{\nu'}^l$ and therefore from (13b) there exist no \tilde{x}, \tilde{y} s.t. $\tilde{y} \in \rho_\nu^l(\tilde{x}, y(r'))$, implying $r' = \text{end}(\pi)$ which is a contradiction to the assumption. \square

Lemma 7: Let π be a play computed by (25) up to length $\text{end}(\pi)$. Then it holds for all $l \in [1, L]$ and $k < \kappa^l(\text{end}(\tilde{\pi}^l))$ that $\neg \text{Done}^l$.

Proof: We prove the statement by contradiction. Pick any $l \in [1, L]$ and $k < \kappa^l(\text{end}(\tilde{\pi}^l))$ and assume that Done^l is true. First observe that this implies $\text{Done}^{l'}$ for all $l' \in [l, L]$. With this it follows from (25f) that $f^l(k) = h^l$. Now using (25k) this implies that $(\tilde{\gamma}^l(k), x_\downarrow^l(k+1)) \notin \text{dom}(f^l(k))$ and therefore the play would not be able to leave the current context. This is a contradiction to the assumption that $k < \kappa^l(\text{end}(\tilde{\pi}^l))$, what proves the statement. \square

Lemma 8: Let π be a play computed by (25) up to length $\text{end}(\pi)$. Then it holds for all $l \in [0, L]$ and $m \in \text{dom}^+(\tilde{\pi}^l)$ that

$$\tilde{x}_\downarrow^l(m) \in \delta_{\tilde{y}^l(m-1)}^l(\tilde{x}_\downarrow^l(m-1), \tilde{y}^l(m-1)) \text{ and} \quad (44a)$$

$$\tilde{y}^l(m) = f^l(\kappa^l(m) - 1)(\tilde{\gamma}^l(\kappa^l(m) - 1), \tilde{x}_\downarrow^l(m)). \quad (44b)$$

Proof: Recall that $\pi \in \mathcal{G}$ from Prop. 4. Therefore, (44a) follows directly from (40) in Lem. 3 (see App.). We show (44b) by induction.

► $l = 0$:

Recall that (27) holds for $l = 0$. As κ^0 is the identity map, the second line in (27) and (44) is equivalent for $l = 0$.

► $l \rightarrow l+1$:

• Pick $m \in \text{dom}^+(\tilde{\pi}^{l+1})$, $k = \kappa^{l+1}(m)$, $\nu := \tilde{y}^{l+1}(m-1)$ and $\nu' := \tilde{y}^{l+1}(m)$ and recall from Lem. 1 that there exists $r \in \mathbb{N}$ s.t. $r = \kappa_l^{l+1}(m)$ and $\kappa^l(r) = k$, implying (from (7)) that

$$y^{l+1}(k-1) = \nu \neq \nu' = y^{l+1}(k), \quad y^l(k) = \tilde{y}^l(r), \quad (45a)$$

$$\tilde{x}_\downarrow^{l+1}(m) = x_\downarrow^{l+1}(k), \text{ and } \tilde{x}_\downarrow^l(r) = x_\downarrow^l(k).$$

Now it follows from Lem. 3 that $\tilde{p}_\downarrow^l(m-1) \in \mathcal{G}_\nu^l$, hence Lem. 6 holds for $\tilde{p}_\downarrow^l(m-1)$. Now using the first and second line of (36a) in Lem. 6 immediately implies

$$\tilde{p}_\downarrow^l(m-1) \in \psi_\nu^l(\nu'). \quad (45b)$$

• We now show that $\tilde{p}_\downarrow^l(m-1)$ is compliant with $f^l(k-1)$:

As (44) holds for l we know that for all $k' < \text{end}(\pi)$ we have $\neg \text{GotStuck}^{l+1}(k')$ (from Lem. 4). As additionally $\neg \text{Done}^{l+1}(k')$ from Lem. 7, (25f) gives that

$$f^l(k') = f_{y^{l+1}(k')\nu^{l+1}}^l(k') \quad (45c)$$

$$\text{s.t. } \nu^{l+1}(k') = f^{l+1}(k')(\tilde{\gamma}^{l+1}(k'), x_\downarrow^{l+1}(k'+1)). \quad (45d)$$

Now pick s s.t.

$$\forall k', k'' \in [\kappa^l(r-s), \kappa^l(r)-1] . y^{l+1}(k') = y^{l+1}(k'') \wedge \nu^{l+1}(k') = \nu^{l+1}(k''), \quad (45e)$$

with ν^{l+1} as in (45d) and observe that this implies $\kappa^l(r-s) \in [\kappa^{l+1}(m-1), \kappa^{l+1}(m))$. Using (45e) in (25i) therefore gives for all $k' \in [\kappa^l(r-s), \kappa^l(r)-1]$ that

$$f^l(k') = f^l(k-1) = \text{Sol}(G_\nu^l, \{\tilde{\gamma}^l(\kappa^l(r-s))\}, \phi_\nu^l(\nu^{l+1}(k-1))). \quad (45f)$$

As (44) holds for l we can therefore substitute $f^l(\kappa^l(r)-1)$ in (44b) by $f^l(k-1)$ and obtain for all $r' \in [r-s, r-1]$ that

$$\tilde{y}^l(r') = f^l(k-1)(\tilde{\gamma}^l(\kappa^l(r'))-1, \tilde{x}_\downarrow^l(r')). \quad (45g)$$

It furthermore follows from the construction of $\check{\gamma}^l$ in (25o) and \check{p}_\downarrow^l in (16) that

$$\check{p}_\downarrow^l(m-1) = \check{\gamma}^l(\kappa^l(r-s)) \cdot \check{\pi}_\downarrow^l|_{[r-s+1, r]} \quad (45h)$$

Now pick $n = \text{end}(\check{\gamma}^l(\kappa^l(r-s)))$ and observe that $\check{p}_\downarrow^l(m-1)|_{[0, n]} \in \{\check{\gamma}^l(\kappa^l(r-s))\}$. Additionally using (45g) therefore implies that $\check{p}_\downarrow^l(m-1) \in \text{CompliantPlays}(f^l(k-1), \{\check{\gamma}^l(\kappa^l(r-s))\})$ (from (2)). Using (45f) and (24b) from Prop. 3 it follows that

$$\check{p}_\downarrow^l(m-1) \in \overline{\phi_\nu^l(\nu^{l+1}(k-1))}. \quad (45i)$$

• It remains to shown that $\nu^{l+1}(k-1) = \nu' (= \check{y}^{l+1}(m) = y^{l+1}(k))$:

Using the fact that $y^l(k) \in Y_{\nu^l}^l$ it follows from Lem. 6 and (21) that (45b) and (45i) can only be satisfied simultaneously if

$$\check{p}_\downarrow^l(m-1) \in \phi_\nu^l(\nu^{l+1}(k-1)) \quad \text{and} \quad \nu^{l+1}(k-1) = \nu'. \quad (45j)$$

With this observation (44b) immediately follows for $l+1$ from (45d) as $\nu^{l+1}(k-1) = \check{y}^{l+1}(m)$. \square

Lemma 9: Let π be a play computed by (25) up to length $\text{end}(\pi)$ and $[\check{p}]_\pi$ its induced set of local projected play sequences. Then it holds for all $l \in [0, L-1]$ and $m \in \text{dom}^+(\check{\pi}^{l+1})$ that

$$\check{p}_\downarrow^l(m-1) \in \phi_{\check{y}^{l+1}(m-1)}^l(\check{y}^{l+1}(m)) \quad (46a)$$

and for $m = \text{end}(\check{\pi}^{l+1})$ there exists $\nu' \in \text{Post}^{l+1}(\check{y}^{l+1}(m))$ s.t.

$$\check{p}_\downarrow^l(m) \in \overline{\phi_{\check{y}^{l+1}(m)}^l(\nu')}. \quad (46b)$$

Proof: (46a) follows from (45j) in the proof of Lem. 8. We prove (46b):

Pick $l \in [0, L-1]$ and $m = \text{end}(\check{\pi}^{l+1})$ and recall from Lem. 8 that (44) holds for all $k \in \text{dom}^+(\check{\pi}^l)$. Therefore $\neg \text{GotStuck}^{l+1}(k')$ (from Lem. 4) for all $k' < \text{end}(\pi)$. Now we have two cases.

(i) If $\neg \text{Done}^{l+1}(\kappa^{l+1}(m))$, (45d) in the proof of Lem. 8 holds for $k' \in [\kappa^{l+1}(m), \kappa^l(\text{end}(\check{\pi}^l))]$. Following exactly the same reasoning as in (45d)-(45i) we obtain

$$\check{p}_\downarrow^l(m) \in \overline{\phi_{\check{y}^{l+1}(m)}^l(\nu^{l+1}(k-1))}$$

with $\nu^{l+1}(k-1)$ as in (45d), implying (46b).

(ii) If $\text{Done}^{l+1}(\kappa^{l+1}(m))$, it follows from (25f) and (25h) that for $k' \in [\kappa^{l+1}(m), \kappa^l(\text{end}(\check{\pi}^l))]$

$$f^l(k') = h^l(\kappa^{l+1}(m)) = \text{Sol} \left(G_{\check{y}^{l+1}(m)}^l, \{\check{\gamma}^l(\kappa^{l+1}(m))\}, \varphi_{\check{y}^{l+1}(m)}^l, \zeta_\nu^l \right) \quad (47)$$

and from the construction of $\check{\gamma}^l$ and \check{p}_\downarrow^l in (25o) and (16) that $\check{\gamma}^l(\kappa^{l+1}(m)) = [\check{p}_\downarrow^l(m-1)]$. By substituting (47) in (44b) we therefore obtain $\check{p}_\downarrow^l(m) \in \text{CompliantPlays}(h^l(\kappa^{l+1}(m)), [\check{p}_\downarrow^l(m-1)])$ (from (2)). Using (45f) and (24b) from Prop. 3 it follows that $\check{p}_\downarrow^l(m) \in \overline{\phi_{\check{y}^{l+1}(m)}^l(\nu^{l+1}(k-1))}$. Now recall from (21) that $\overline{\phi_{\check{y}^{l+1}(m)}^l} \subseteq \overline{\phi_{\check{y}^{l+1}(m)}^l(\nu^{l+1}(k-1))}$, what proves the statement. \square

Lemma 10: Let π be a maximal and environment admissible play computed by (25) s.t. (33) holds. Then it holds that

$$(\exists k \in \text{dom}(\pi), l \in [0, L] . \text{Done}^l(k)) \Rightarrow (\exists k' \in \text{dom}(\pi), k' \geq k . \text{Done}^0(k')).$$

Proof: Pick $k \in \text{dom}(\pi), l \in [0, L]$ s.t. $\text{Done}^l(k)$ and assume $l > 0$ as for $l = 0$ the statement follows trivially. Giving $\text{Done}^l(k)$, (25l) implies $\neg \text{GotStuck}^l(k)$ and (25f) implies $f^{l-1}(k) = h^{l-1}(k)$. Giving $\neg \text{UnRealizable}^{l-1}(k)$ and $\neg \text{GotStuck}^l(k)$, (26) implies $\neg \text{GotStuck}^{l-1}(k)$ and therefore (from (25l)) either $\text{Done}^{l-1}(k)$ or there exists a next step according to $h^{l-1}(k)$. Assume the latter is true. Recall from (25o) that $h^{l-1}(k)$ is an assume admissible winning strategy for the game $(G_{y^l(k)}^l, \{\check{\gamma}^{l-1}(\lceil \kappa^{l-1} \rceil)\}, \varphi_{y^l(k)}^l)$ and from (18) that $\varphi_{y^l(k)}^l$ only contains finite strings. If the environment plays admissible, we therefore eventually obtain $\text{Done}^{l-1}(k')$ with $k < k' < \infty$. Applying this reasoning iteratively, eventually leads to $\text{Done}^0(k'')$ where the time between k and k'' is ensured to be finite. \square

Lemma 11: Let π be a maximal and environment admissible play computed by (25) s.t. (33) holds. Then it holds that

$$(\forall k \in \text{dom}(\pi), l \in [0, L] . \neg \text{Done}^l(k)) \Leftrightarrow (\forall l \in [0, L] . |\check{\pi}^l| = \infty) \Leftrightarrow (|\pi| = \infty).$$

Proof: We show this proof in two steps.

► Show $(\forall k \in \text{dom}(\pi), l \in [0, L] . \neg \text{Done}^l(k)) \Leftrightarrow (|\pi| = \infty)$:

Using (33) in (29) of Thm. 1 gives

$$\exists l \in [0, L] . \forall k \in \text{dom}(\pi) . \neg \text{Done}^l(k) \Leftrightarrow |\pi| = \infty, \quad (48)$$

immediately implying the “ \Rightarrow ” part of the statement. Now we prove the “ \Leftarrow ” part by contradiction. Assume that there exists $l \in [0, L], k \in \text{dom}(\pi)$ s.t. $\text{Done}^l(k)$. Then Lem. 10 implies $\text{Done}^0(k')$. Using (from (25k)) this implies $\text{Done}^l(k')$ for all $l \in [0, L]$, which gives a contradiction as the left side of (48) holds from $(|\pi| = \infty)$.

► Show $\forall l \in [0, L] . |\check{\pi}^l| = \infty \Leftrightarrow |\pi| = \infty$:

First observe that “ \Rightarrow ” trivially holds as $\check{\pi}^0 = \pi$. We prove “ \Leftarrow ” by contradiction. Assume there exists $l \in [0, L]$ s.t.

$|\tilde{\pi}^l| < \infty$, i.e., with $k = \text{end}(\tilde{\pi}^l)$ we have $(\gamma^l(k), x_{\downarrow}^l(k+1)) \notin \text{dom}(f^l(k))$. Now recall from the first part of this proof that $|\pi| = \infty$ implies $\neg \text{Done}^{l'}(k)$ for all $l' \in [0, L]$ and (33) implies $\neg \text{UnRealizable}^{l+1}(k)$. Then it follows from Lem. 4 that $\text{GotStuck}^{l+1}(k)$. With this $\text{GotStuck}^l(k)$ (from (25f)) and therefore eventually $\text{GotStuck}^0(k)$, which implies $|\pi| < \infty$ with $\text{end}(\pi) = k$, which is a contradiction to the assumption. \square

REFERENCES

- M. Abadi and L. Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82(2):253 – 284, 1991.
- R. Alur, S. La Torre, and P. Madhusudan. Modular strategies for recursive game graphs. In *TACAS*, volume 2619 of *LNCS*, pages 363–378. 2003.
- R. Bloem, R. Ehlers, S. Jacobs, and R. Könighofer. How to handle assumptions in synthesis. In *SYNT 2014, Vienna, Austria*, pages 34–50, 2014.
- R. Brenguier, J.-F. Raskin, and M. Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS*, volume 23, pages 1–10, 2014.
- R. Brenguier, J. Raskin, and O. Sankur. Assume-admissible synthesis. *CoRR*, 2015.
- P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL 77*, pages 238–252. ACM, 1977.
- I. De Crescenzo and S. La Torre. Modular synthesis with open components. In *Reachability Problems*, volume 8169 of *LNCS*, pages 96–108. 2013.
- R. Ehlers and B. Finkbeiner. Reactive safety. In *GandALF, EPTCS 54*, pages 178–191, 2011.
- E. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS*, pages 368–377, 1991.
- K. Erol, J. A. Hendler, and D. S. Nau. Semantics for hierarchical task-network planning. Technical report, University of Maryland, 1995.
- C. Finucane, G. Jing, and H. Kress-Gazit. LTLMoP: Experimenting with language, temporal logic and robot control. In *IROS*, pages 1988–1993, 2010.
- A. Girard and G. J. Pappas. Hierarchical control system design using approximate simulation. *Automatica*, 45(2):566 – 571, 2009.
- T. A. Henzinger, R. Majumdar, F. Mang, and J.-F. Raskin. Abstract interpretation of game properties. In *Static Analysis*, LNCS 1824, pages 220–239. 2000.
- D. Hess, M. Althoff, and T. Sattel. Formal verification of maneuver automata for parameterized motion primitives. In *IROS*, pages 1474–1481, Sept 2014.
- L. Kaelbling and T. Lozano-Perez. Hierarchical task and motion planning in the now. In *ICRA*, pages 1470–1477, May 2011.
- M. Kloetzer and C. Belta. Dealing with nondeterminism in symbolic control. In *HSCC*, volume 4981 of *LNCS*, pages 287–300. 2008.
- T. Koo and S. Sastry. Bisimulation based hierarchical system architecture for single-agent multi-modal systems. In *HSCC*, volume 2289 of *LNCS*, pages 281–293. 2002.
- N. Kruger, J. Piater, F. Worgotter, C. Geib, R. Petrick, M. Steedman, T. Asfour, D. Kraft, B. Hommel, A. Agostini, et al. A formal definition of object-action complexes and examples at different levels of the processing hierarchy. *Computer and Information Science*, pages 1–39, 2009.
- O. Kupferman and M. Vardi. Model checking of safety properties. *Formal Methods in System Design*, 19(3):291–314, 2001.
- O. Kupferman and S. Weiner. Environment-friendly safety. In *HVC 2012*, volume 7857 of *LNCS*, pages 227–242. Springer, 2012.
- J. Mazo, Manuel, A. Davitian, and P. Tabuada. PESSOA: A tool for embedded controller synthesis. In *CAV*, volume 6174 of *LNCS*, pages 566–569. 2010.
- G. Pappas, G. Lafferriere, and S. Sastry. Hierarchically consistent control systems. *Trans. on Automatic Control*, 45(6): 1144–1160, Jun 2000.
- J. Raisch and T. Moor. Hierarchical hybrid control synthesis and its application to a multiproduct batch plant. In *Control and Observer Design for Nonlinear Finite and Infinite Dimensional Systems*, volume 322 of *LNCS*, pages 199–216. 2005.
- T. Reps, S. Horwitz, and S. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *POPL 95*, pages 49–61. ACM, 1995.
- K. Schmidt, T. Moor, and S. Perk. Nonblocking hierarchical control of decentralized discrete event systems. *IEEE Transactions on Automatic Control*, 53(10):2252–2265, 2008.
- S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel. Combined task and motion planning through an extensible planner-independent interface layer. In *ICRA*, pages 639–646, May 2014.
- S. Stock, M. Mansouri, F. Pecora, and J. Hertzberg. Hierarchical hybrid planning in a mobile service robot. In *KI 2015: Advances in Artificial Intelligence*, pages 309–315. 2015.
- P. Tabuada. *Verification and Control of Hybrid Systems - A Symbolic Approach*, volume 1. Springer, 2009.

- C. I. Vasile and C. Belta. Reactive sampling-based temporal logic path planning. In *ICRA*, pages 4310–4315, 2014.
- I. Walukiewicz. Pushdown processes: Games and model checking. In *CAV 96: Computer-Aided Verification*, LNCS 1102, pages 62–74, 1996.
- E. Wolff, U. Topcu, and R. Murray. Optimal control of non-deterministic systems for a computationally efficient fragment of temporal logic. In *CDC*, pages 3197–3204, 2013.
- K. W. Wong, C. Finucane, and H. Kress-Gazit. Provably-correct robot control with ltlmop, ompl and ros. In *IROS*, pages 2073–2073, Nov 2013.
- T. Wongpiromsarn, U. Topcu, and R. M. Murray. Automatic synthesis of robust embedded control software. In *AAAI Spring Symposium: Embedded Reasoning*, 2010.
- T. Wongpiromsarn, U. Topcu, N. Ozay, H. Xu, and R. M. Murray. TuLiP: A software toolbox for receding horizon temporal logic planning. In *HSCC*, pages 313–314, 2011.
- T. Wongpiromsarn, U. Topcu, and R. Murray. Receding horizon temporal logic planning. *Transactions on Automatic Control*, 57(11):2817–2830, 2012.
- W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1-2):135–183, 1998.